Santa Fe Institute Working Paper 15-10-XXX
arxiv.org:1510.XXXX [physics.gen-ph]

A Closed-Form Shave from Occam's Quantum Razor: Exact Results for Quantum Compression

Paul M. Riechers,* John R. Mahoney,† Cina Aghamohammadi,‡ and James P. Crutchfield§
Complexity Sciences Center and Department of Physics
University of California at Davis
One Shields Avenue, Davis, CA 95616
(Dated: October 29, 2015)

The causal structure of a stochastic process can be more efficiently transmitted via a quantum channel than a classical one, an advantage that increases with codeword length. While previously difficult to compute, we express the quantum advantage in closed form using spectral decomposition, leading to direct computation of the quantum communication cost at all encoding lengths, including infinite. This makes clear how finite-codeword compression is controlled by the classical process' cryptic order and allows us to analyze structure within the length-asymptotic regime of infinite-cryptic order (and infinite Markov order) processes.

PACS numbers: 03.67.Hk 03.67.-a 03.65.-w 03.65.Db
Keywords: epsilon-machine, information compression, quantum state overlap, crypticity, spectral decomposition

arXiv:1510.08186v1 [quant-ph] 28 Oct 2015

I. INTRODUCTION

Defining and quantifying "pattern" and "structure" have been active endeavors for decades [1–4]. They are especially important and challenging when it comes to the complex patterns spontaneously generated by nonlinear dynamical systems and hidden stochastic processes. These studies seek to answer questions such as, "how unpredictable is this pattern?" and "how much memory is needed to accurately predict its next element?" Computational mechanics [5, 6], an extension of statistical mechanics, was established to answer these very questions. In much of computational mechanics, the emphasis is on stationary stochastic processes. Here, we explore how their internal structure generates patterns.

The question of a process's internal structure naturally arises for two observers, Alice and Bob, who wish to efficiently synchronize their predictions of a given process over a classical communication channel. What is the minimal amount of information they that must communicate? The answer is given by the process' statistical complexity C_μ [1].

A closely related question immediately suggests itself: is it more efficient to synchronize via a quantum communication channel? Extending early answers [7, 8], we recently introduced a sequence of constructions (q-machines) that offer substantially improved quantum compression [9]. Each codeword length L yields a quantum communication rate $C_q(L)$, where $C_q(L) \leq C_\mu$. Moreover, we showed that maximum compression $C_q(\infty) = C_q(k)$ is achieved at a codeword length called the *cryptic order* k [10]—a recently discovered classical, topological property that is a cousin to the more familiar stochastic process Markov order.

The following develops the analytical underpinnings for the q-machine construction and for computing the quantum communication costs $C_q(L)$ [9]. We present a closed-form expression for the overlaps between quantum signal states—overlaps that lead to quantum compressibility. Overlaps are expressed in terms of the spectrum (and related projection operators) of the quantum pairwise-merger machine, introduced here. This leads us to a decomposition of the quantum compression into two qualitatively distinct parts: a finite-horizon contribution at any cryptic order and a persistent infinite-horizon contribution that arises only for infinite-cryptic processes.

An even greater advantage obtains when using a surrogate for the quantum density matrix—a fixed-size Gram matrix, which has the same spectrum. Critically, the surrogate matrix is linear in the overlaps and is straightforward to calculate. Using it, the quantum communication costs $C_q(L)$ (and $C_q(\infty)$) are calculable from a closed-form expression, in some cases analytically and, in all cases, with an unprecedented level of numerical accuracy and efficiency. Moreover, the generic behavior of $C_q(L)$ for large L and infinite cryptic order can be probed. One lesson, for example, is that optimal quantum compression is achieved asymptotically, but with exponentially diminishing returns. And, we discover and explain why the quantum advantage oscillates with increasing codeword length: the oscillations reflect the structure of a process' crypticity—how its internal state information, hidden even upon all subsequent observations, reveals itself only slowly with the additional information of past

* pmriechers@ucdavis.edu
† jrmahoney@ucdavis.edu
‡ caghamohammadi@ucdavis.edu
§ chaos@ucdavis.edu

observations provided with increasing codeword length.

The next section reviews the causal-state representation that computational mechanics employs for stochastic processes—the ϵ-machine—and recalls the new representation—the q-machine—appropriate to quantum communication. We then show how to build quantum codewords out of a classical process' causal states and observed sequences and how to monitor the codewords' quantum overlap. Central to determining the proper overlap structure, we introduce the quantum pairwise-merger machine (QPMM), detailing an algorithm for its construction and illustrating its use via several example processes. Next, we turn to measuring quantum compressibility via the von Neumann entropy of the overlap density matrix. Spectral decomposition of the QPMM transition dynamic, though, leads to closed-form expressions for the elements of the Gram matrix, which yields the quantum coding cost at any length. We explore several example processes with these tools to illustrate a number of features when quantum compressing structured classical processes. These include finite and infinite cryptic-order processes and the universal behavior of quantum compression in the asymptotic codeword-length limit. We conclude with a broader discussion of the application of these results and brief comments on how they affect our view of the mechanisms by which classical and quantum processes generate the patterns we see in the physical world.

II. TWO REPRESENTATIONS OF A PROCESS

The objects of interest are discrete-valued, stationary, stochastic processes generated by finite hidden Markov models (HMMs). In particular, we consider edge-output HMMs (or Mealy HMMs) where the observed symbol is generated upon transition between states. Rather than focus on generating models, more prosaically we can also think of a *process* consisting of a bi-infinite sequence $X_{-\infty:\infty} = \ldots X_{-2}X_{-1}X_0X_1X_2\ldots$ of random variables X_t that take on one or another value in a discrete alphabet: $x_t \in \mathcal{A}$. A process' *language* is that set of words $w = x_0 \ldots x_{L-1}$ of any length L generated with positive probability. We consider two representations of a given process, first a canonical classical representation and then a newer quantum representation.

A. ϵ-Machine

While a given process generally has many alternative HMM representations, there exists a unique, canonical form—the process's ϵ-*machine* [1]. An equivalence relation applied to the random variable sequence $X_{-\infty:\infty}$ defines the process' causal states, which encapsulate all about individual pasts relevant for predicting the future. Said another way, causal states are the minimal sufficient statistic of the past $X_{\infty:0}$ for predicting the future $X_{0:\infty}$. (Note that we use array indexing that is left inclusive, but right exclusive.)

Definition 1. *A process'* ϵ-machine $\mathcal{M}$ *is the tuple* $\left(\boldsymbol{\mathcal{S}}, \mathcal{A}, \{T^{(x)}\}_{x\in\mathcal{A}}, \boldsymbol{\pi}\right)$.

$\boldsymbol{\mathcal{S}}$ is the set $\{\sigma_0, \sigma_1, \ldots\}$ of the process' causal states, $\mathcal{A}$ is the set of possible symbol outputs x, $\left\{T^{(x)} : T^{(x)}_{i,j} = \Pr(\sigma_j, x|\sigma_i)\right\}_{x\in\mathcal{A}}$ are the labeled transition matrices, and $\boldsymbol{\pi}$ the stationary distribution over states. The probability that a word w is generated by the ϵ-machine is given in terms of the labeled transition matrices and the initial state distribution:

$$\Pr(w) = \boldsymbol{\pi} \prod_{i=0}^{L-1} T^{(x_i)} \mathbb{1},$$

where $\mathbb{1} = [1, \ldots, 1]^\top$. When these probabilities are constructed to agree with those of the words in a given process language, the ϵ-machine is said to *generate* or *represent* that process.

The ensemble temporal evolution of internal state probability $\boldsymbol{\mu} = \left(\mu_0, \ldots, \mu_{|\boldsymbol{\mathcal{S}}|-1}\right)$, with $\mu_i = \Pr(\sigma_i)$, is given by:

$$\boldsymbol{\mu}(t+1) = \boldsymbol{\mu}(t)T \ , \tag{1}$$

where the net transition matrix T is the sum over all output symbols:

$$T := \sum_{x\in\mathcal{A}} T^{(x)} \ . \tag{2}$$

Transition probabilities are normalized. That is, the transition matrix T is *row-stochastic*:

$$\sum_{j=1}^{|\boldsymbol{\mathcal{S}}|} T_{i,j} = \sum_{j=1}^{|\boldsymbol{\mathcal{S}}|} \sum_{x\in\mathcal{A}} \Pr(\sigma_j, x|\sigma_i) = 1 \ .$$

Its component matrices $T^{(x)}_{ij}$ are said to be *substochastic*. Under suitable conditions on the transition matrix, $\lim_{t\to\infty} \boldsymbol{\mu}(t) = \boldsymbol{\pi}$.

Unifilarity, an inherent property of ϵ-machines, means that for each state σ_i, each symbol x may lead to at most one successor state σ_j [11]. In terms of the labeled transition matrices, for each row i and each symbol x the row $T^{(x)}_{ij}$ has at most one nonzero entry. This property falls naturally out of the equivalence relation used to define causal states [1]. We also will have occasion to speak of a

counifilar HMM, which is the analogous requirement of unique labeling on transitions coming *into* each state.

One of the most important complexity measures for a process is its statistical complexity C_μ [1]. Used in a variety of contexts, it quantifies a process' minimal description size.

Definition 2. *The* statistical complexity C_μ *of a process is the Shannon entropy of the stationary distribution over its ϵ-machine's causal states:*

$$\begin{aligned} C_\mu &= H[\boldsymbol{\pi}] \\ &= -\sum_{i=1}^{|\boldsymbol{\mathcal{S}}|} \pi_i \log \pi_i \ . \end{aligned}$$

The statistical complexity has several operational meanings. For example, it is the average amount of information one gains on determining in which causal state a process is. It is also the minimal amount of information that must be stored from the past to optimally predict the future. Most pertinent to our purposes here, it also quantifies the communication cost of synchronizing two predicting agents through a classical channel [9].

B. q-Machine

The q-machine is a representation of a classical process that makes use of quantum mechanics. Introduced in Ref. [9], it offers the process' most complete quantum compression known so far.

A process' q-machine is constructed by first selecting a *codeword length* L. The q-machine (at L) consists of a set $\{|\eta_i(L)\rangle\}_{i=1}^{|\boldsymbol{\mathcal{S}}|}$ of pure quantum *signal states* that are in one-to-one correspondence with the classical causal states $\sigma_i \in \boldsymbol{\mathcal{S}}$. Each signal state $|\eta_i(L)\rangle$ encodes the set of length-L words $\{w : \Pr(w|\sigma_i) > 0\}$ that may follow causal state σ_i, as well as the corresponding conditional probability:

$$|\eta_i(L)\rangle \equiv \sum_{w \in \mathcal{A}^L} \sum_{\sigma_j \in \boldsymbol{\mathcal{S}}} \sqrt{\Pr(w, \sigma_j|\sigma_i)}\ |w\rangle\, |\sigma_j\rangle \ , \quad (3)$$

where $\left\{ |w\rangle \right\}_{w \in \mathcal{A}^L}$ denotes an orthonormal basis in the "word" Hilbert space with one dimension for each possible word w of length L. Similarly, $\left\{ |\sigma_j\rangle \right\}_{j=1}^{|\boldsymbol{\mathcal{S}}|}$ denotes an orthonormal basis in the "state" Hilbert space with one dimension for each classical causal state. The ensemble of length-L quantum signal states is then described by the density matrix:

$$\rho(L) = \sum_{i=1}^{|\boldsymbol{\mathcal{S}}|} \pi_i\, |\eta_i(L)\rangle\, \langle \eta_i(L)| \ . \quad (4)$$

The ensemble's von Neumann entropy (VNE) is defined in terms of its density matrix: $S(\rho) = -\mathrm{tr}[\rho \log(\rho)]$, where $\mathrm{tr}[\cdot]$ is the trace of its argument. Paralleling the classical statistical complexity, the quantity:

$$\begin{aligned} C_q(L) &\equiv S(\rho(L)) \\ &= -\mathrm{tr}\big[\rho(L) \log(\rho(L))\big] \end{aligned} \quad (5)$$

has the analogous operational meaning of the communication cost over a quantum channel. This is of interest since it is generically smaller [9]: $C_q(L) \leq C_\mu$. In fact, $C_\mu = C_q$ if and only if the process' ϵ-machine is counifilar—there are no states with (at least) two similarly labeled incoming edges [7]. Notably, processes with counifilar ϵ-machines are a vanishing proportion of all processes, as we increase state size [12]. The consequence is that almost all classical processes can be more compactly represented using quantum mechanics. This presents an opportunity to use quantum encoding to more efficiently represent a process.

Actually computing the quantum cost $C_q(L)$ is challenging, however. Circumventing this difficulty is one of the main motivations for the following developments. The eventual results, though, allow us to move beyond this practical concern to a deeper appreciation of quantum structural complexity.

III. QUANTUM OVERLAPS

Reference [9] showed that the compression advantage $C_\mu - C_q(L)$ is determined by quantum overlaps between signal states in the q-machine. Moreover, these overlaps are determined by words whose state paths merge.

To illustrate, we compute some overlaps for the $(R$–$k)$-Golden Mean Process, showing how they depend on L. (See Fig. 1 for its ϵ-machine state transition diagram.) This process was designed to have a tuneable Markov order R and cryptic order k; here we choose $R = 4$ and $k = 3$. (Refer to Ref. [12] for more on this process and detailed discussion of Markov and cryptic orders.)

At length $L = 0$, each signal state is simply the basis state corresponding to its causal state: $|\eta_i(0)\rangle = |\sigma_i\rangle$. Since the ϵ-machine is minimal, there are no overlaps in the state vectors.

At length $L = 1$ codewords, we find the first nontrivial overlap. This corresponds to paths $A \xrightarrow{1} A$ and $G \xrightarrow{1} A$

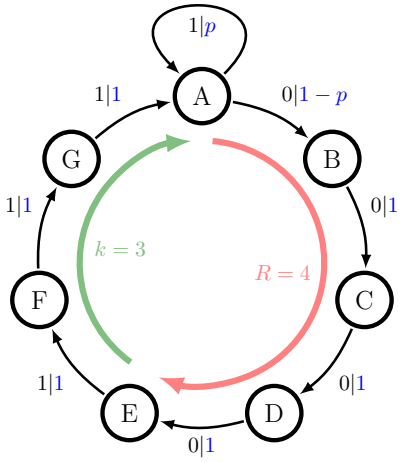

FIG. 1. ϵ-Machine for the (4–3)-Golden Mean Process: The cycle's red part indicates the "Markov" portion and the green, the "cryptic" portion. The time scales R and k are tuned by changing the lengths of these two parts. Edges labeled $x|p$ denote taking the state-to-state transition with probability p while emitting symbol $x \in \mathcal{A}$.

merging at state A and we have:

$$|\eta_A(1)\rangle = \sqrt{p}\,|1A\rangle + \sqrt{1-p}\,|0B\rangle \text{ and}$$
$$|\eta_G(1)\rangle = |1A\rangle \ .$$

This yields the overlap:

$$\langle \eta_A(1)|\eta_G(1)\rangle = \sqrt{p}.$$

Going on to length $L = 2$ codewords, more overlaps arise from mergings of more state paths. The three quantum signal states:

$$|\eta_A(2)\rangle = p\,|11A\rangle + \sqrt{p(1-p)}\,|10B\rangle + \sqrt{(1-p)}\,|00C\rangle \ ,$$
$$|\eta_F(2)\rangle = |11A\rangle \ , \text{ and}$$
$$|\eta_G(2)\rangle = \sqrt{p}\,|11A\rangle + \sqrt{1-p}\,|10B\rangle$$

interact to yield the overlaps:

$$\langle \eta_A(2)|\eta_F(2)\rangle = p \ ,$$
$$\langle \eta_F(2)|\eta_G(2)\rangle = \sqrt{p} \ , \text{ and}$$
$$\langle \eta_A(2)|\eta_G(2)\rangle = p\sqrt{p} + (1-p)\sqrt{p} = \sqrt{p} \ .$$

The overlaps between (A, F) and (F, G) are new. The (A, G) overlap has the same value as that for (F, G), however its computation at $L = 2$ involved two terms instead of one. This is because no new merger has occurred; the $L = 1$ merger, affected by symbol 1, was simply propagated forward along two different state paths having prefix 1. The redundant paths are: $A \xrightarrow{10} B$ overlaps $G \xrightarrow{10} B$ and $A \xrightarrow{11} A$ overlaps $G \xrightarrow{11} A$. A naive computation of overlaps must contend with this type of redundancy.

IV. QUANTUM PAIRWISE-MERGER MACHINE

To calculate signal-state overlaps, we introduce the quantum pairwise-merger machine, a transient graph structure that efficiently encapsulates the structure of state paths. As we saw in the example, computation of overlaps amounts to tracking state-path mergers. It is important that we do this in a systematic manner to avoid redundancies. The new machine does just this.

We begin by first constructing the *pairwise-merger machine* (PMM), previously introduced to compute overlaps [9]. There, probabilities were computed for each word found by scanning through the PMM. This method significantly reduced the number of words from the typically exponential large number in a process' language and also gave a stopping criterion for PMMs with cycles. This was a vast improvement over naive constructions of the signal-state ensemble (just illustrated) and over von Neumann entropy calculation via diagonalization of an ever-growing state space.

Appropriately weighting PMM transitions yields the *quantum PMM* (QPMM), which then not only captures which states merge given which words, but also the contribution each merger makes to a quantum overlap. The QPMM has one obvious advantage over the PMM. The *particular* word that produces an overlap is ultimately unimportant; only the amount of overlap generated is important. Therefore, summing over symbols in the QPMM to obtain its internal state transitions removes this combinatorial factor. Appreciating other significant advantages to this matrix-based approach requires more development first.

To build the QPMM from a given process' ϵ-machine:

1. Construct the set of (unordered) pairs of (distinct) ϵ-machine states: (σ_i, σ_j). We call these "pair-states". To this set, add a special state called SINK (short for "sink of synchronization") which is the terminal state.
2. For each pair-state (σ_i, σ_j) and each symbol $x \in \mathcal{A}$, there are three cases to address:
   (a) If at least one of the two ϵ-machine states σ_i or σ_j has no outgoing transition on symbol x, then do nothing.
   (b) If both ϵ-machine states σ_i and σ_j have a transition on symbol x to the same state σ_m, then connect pair-state (σ_i, σ_j) to SINK with an edge labeled x. This represents a merger.

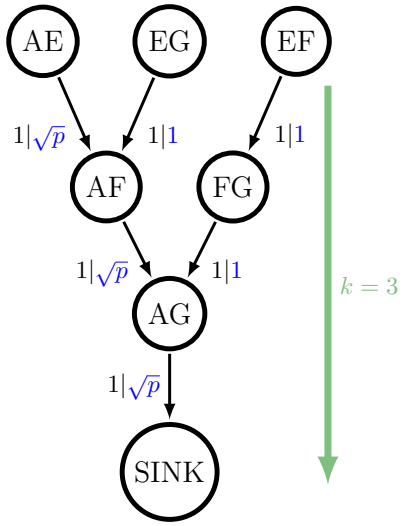

FIG. 2. QPMM for the (4–3)-Golden Mean Process. Its depth is related to the cryptic order k.

(c) If both ϵ-machine states σ_i and σ_j have a transition on symbol x to two distinct ϵ-machine states σ_m and σ_n where $m \neq n$, then connect pair-state (σ_i, σ_j) to pair-state (σ_m, σ_n) with an edge labeled x. (There are no further restrictions on m and n.)

3. Remove all edges that are not part of a path that leads to SINK.

4. Remove all pair-states that do not have a path to SINK.

This is the PMM. Now, add information about transition probabilities to this topological structure to obtain the QPMM:

5. For each pair-state (σ_i, σ_j) in the PMM, add to each outgoing edge the weight $\sqrt{\Pr(x|\sigma_i)\Pr(x|\sigma_j)}$, where x is the symbol associated with that edge. Note that two states in QPMM may be connected with multiple edges (for different symbols).

Returning to our example, Fig. 2 gives the QPMM for the (4–3)-Golden Mean Process. Using it, we can easily determine the length at which a contribution is made to a given overlap. We consider codeword lengths $L = 1, 2, \ldots$ by walking up the QPMM from SINK. For example, pair (A, G) receives a contribution of $\sqrt{p}$ at $L = 1$. Furthermore, (A, G) receives no additional contributions at larger L. Pairs (A, F) and (F, G), though, receive contributions $p = \sqrt{p} \times \sqrt{p}$ and $\sqrt{p} = \sqrt{p} \times 1$ at $L = 2$, respectively.

The QPMM is *not* a HMM, since the edge weights do not yield a stochastic matrix. However, like a HMM, we can consider its labeled "transition" matrices $\{\zeta^{(x)}\}$, $x \in \mathcal{A}$. Just as for their classical ϵ-machine counterparts, we index these matrices such that $\zeta^{(x)}_{u,v}$ indicates the edge going from pair-state u to pair-state v. Since the overlap contribution, and not the inducing word, is of interest, the important object is simply the resulting state-to-state substochastic matrix $\zeta = \sum_{x \in \mathcal{A}} \zeta^{(x)}$. The matrix ζ is the heart of our closed-form expressions for quantum coding costs, which follow shortly. As we noted above, it is this step that greatly reduces the combinatorial growth of paths that would otherwise make the calculations unwieldy.

To be explicit, our (4–3)-Golden Mean Process has:

$$\zeta = \begin{array}{c} \\ AE \\ EG \\ EF \\ AF \\ FG \\ AG \\ \text{SINK} \end{array} \begin{array}{c} \begin{array}{ccccccc} AE & EG & EF & AF & FG & AG & \text{SINK} \end{array} \\ \begin{pmatrix} 0 & 0 & 0 & \sqrt{p} & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \sqrt{p} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \sqrt{p} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \end{array} .$$

V. OVERLAPS FROM THE QPMM

As we saw in the example, overlaps accumulate contributions as "probability amplitude" is pushed through the QPMM down to the SINK. Each successive overlap augmentation can thus be expressed in terms of the next iterate of ζ:

$$\begin{aligned} &\langle \eta_i(L) | \eta_j(L) \rangle - \langle \eta_i(L-1) | \eta_j(L-1) \rangle \\ &\qquad = \langle (\sigma_i, \sigma_j) | \zeta^L | \text{SINK} \rangle \ . \end{aligned}$$

The general expression for quantum overlaps follows immediately:

$$\langle \eta_i(L) | \eta_j(L) \rangle = \langle (\sigma_i, \sigma_j) | \sum_{n=0}^{L} \zeta^n | \text{SINK} \rangle \ . \qquad (6)$$

This form makes clear the cumulative nature of quantum overlaps and the fact that overlap contributions are not "labeled".

Note that there are two trivial overlap types. Self-overlaps are always 1; this follows from Eq. (6) since $\langle (\sigma_i, \sigma_i) | = \langle \text{SINK} |$. Overlaps with no corresponding pair-state in the QPMM are defined to be zero for all L.

Now, we show that there are two behaviors that contribute to overlaps: a finite-horizon component and an infinite-horizon component. Some processes have only

one type or the other, while many have both. We start with the familiar (R–k)-GM, which has only the finite-horizon behavior.

A. Finite Horizon: (R–k)-Golden Mean Process

Let's use the general expression Eq. (6) to compute the Hermitian positive semidefinite *overlap matrices* $A_L A_L^\dagger$ with components $(A_L A_L^\dagger)_{i,j} = \langle \eta_i(L) | \eta_j(L) \rangle$ for lengths $L = 1, 2, 3, 4$ for the (R–k)-Golden Mean Process. We highlight in blue the matrix elements that have changed from the previous length. All overlaps begin with the identity matrix, here I_7 as we have seven states in the ϵ-machine (Fig. 1). Then, at $L = 1$ we have one overlap. The overlap matrix, with elements $\langle \eta_i(1) | \eta_j(1) \rangle$, is:

$$A_1 A_1^\dagger = \begin{array}{c} \\ A \\ B \\ C \\ D \\ E \\ F \\ G \end{array} \begin{array}{c} \begin{array}{ccccccc} A & B & C & D & E & F & G \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & \sqrt{p} \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ \sqrt{p} & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{array} .$$

Next, for $L = 2$ we find two new overlaps. The overlap matrix, with elements $\langle \eta_i(2) | \eta_j(2) \rangle$, is:

$$A_2 A_2^\dagger = \begin{array}{c} \\ A \\ B \\ C \\ D \\ E \\ F \\ G \end{array} \begin{array}{c} \begin{array}{ccccccc} A & B & C & D & E & F & G \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & p & \sqrt{p} \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ p & 0 & 0 & 0 & 0 & 1 & \sqrt{p} \\ \sqrt{p} & 0 & 0 & 0 & 0 & \sqrt{p} & 1 \end{pmatrix} \end{array} .$$

For $L = 3$, there are three new overlaps. The overlap matrix, with elements $\langle \eta_i(3) | \eta_j(3) \rangle$, is:

$$A_3 A_3^\dagger = \begin{array}{c} \\ A \\ B \\ C \\ D \\ E \\ F \\ G \end{array} \begin{array}{c} \begin{array}{ccccccc} A & B & C & D & E & F & G \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 & \sqrt{p}^3 & p & \sqrt{p} \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ \sqrt{p}^3 & 0 & 0 & 0 & 1 & \sqrt{p} & p \\ p & 0 & 0 & 0 & \sqrt{p} & 1 & \sqrt{p} \\ \sqrt{p} & 0 & 0 & 0 & p & \sqrt{p} & 1 \end{pmatrix} \end{array} .$$

Finally, for $L = 4$, we find the same matrix as $L = 3$: $\langle \eta_i(4) | \eta_j(4) \rangle = \langle \eta_i(3) | \eta_j(3) \rangle$ for all i and j. And, in fact, this is true for all $L \geq 3$. Therefore, all overlap information has been uncovered at codeword length $L = 3$.

Looking at the QPMM in Fig. 2, we recognize that the saturation of the overlap matrix corresponds to the finite *depth* d of the directed graph—the longest state-path through the QPMM that ends in the SINK state. Equivalently, the depth corresponds to the nilpotency of ζ:

$$d = \min\{n \in \mathbb{N} : \zeta^n = 0\} . \tag{7}$$

Note that the $(4-3)$-Golden Mean Process QPMM is a tree of depth 4.

Whenever the QPMM is a tree or, more generally, a directed-acyclic graph (DAG), the overlaps will similarly have a finite-length horizon equal to the depth d. The nilpotency of ζ for finite-depth DAGs allows for a truncated form of the general overlap expression Eq. (6):

$$\langle \eta_i(L) | \eta_j(L) \rangle = \langle (\sigma_i, \sigma_j) | \sum_{n=0}^{\min(L, d-1)} \zeta^n \, | \text{SINK} \rangle . \tag{8}$$

This form is clearly advantageous for any process whose QPMM is a finite DAG. Naturally then, we are led to ask: What property of a process leads to a finite DAG? To answer this question, we reconsider how overlap is accumulated via the merging of state-paths.

Paths through the QPMM represent causal-state-path mergers. To make this more precise, we introduce the concept of an L-merge, which is most intuitively understood through Fig. 3:

Definition 3. *An* L-merge *consists of a length-*L *word* w *and two state paths each of length* $L+1$ *that each generate the word* w *and end at the same state exactly at length* $L+1$. *We denote the word* $w = (x_0, \ldots, x_{L-1})$ *and state paths* $(a_0, \ldots, a_{L-1}, F)$ *and* $(b_0, \ldots, b_{L-1}, F)$ *where states* $a_i \neq b_i$, *for all* $i \in [0, L-1]$ *and, trivially,* $F = F$,

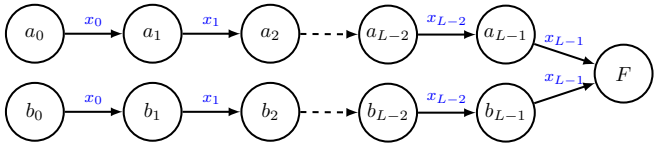

FIG. 3. L-merge: Two causal-state paths—$(a_0, \ldots, a_{L-1}, F)$ and $(b_0, \ldots, b_{L-1}, F)$ where states $a_i \neq b_i$, for all $i \in [0, L-1]$—generate the same word $w = x_0 x_1 \ldots x_{L-1}$ and merge only on the last output symbol x_{L-1} into a common final state F.

the final state in which the paths end.

Immediately, we see that every labeled path of length-L through the QPMM that ends in SINK is an L-merge.

Such causal-state-path merging not only contributes to quantum overlap, but also contributes to a process' cryptticity $\mathrm{H}[\mathcal{S}_0|X_{0:\infty}]$, which is accumulated at all lengths up to the *cryptic order* [13]:

Definition 4. *The* cryptic order k *of a process is the minimum length* L *for which* $\mathrm{H}[\mathcal{S}_L|X_{0:\infty}] = 0$.

A process' *Markov order* is:

$$R = \min\{L : \mathrm{H}[\mathcal{S}_L|X_{0:L}] = 0\} \ .$$

A more familiar length-scale characterizing historical dependence, R depends on both path merging and path termination due to disallowed transitions. The cryptic order, in contrast, effectively ignores the termination events and is therefore upper-bounded by the Markov order: $k \leq R$. This bound is also easy to see given the extra conditional variable $X_{L:\infty}$ in cryptticity ($X_{0:\infty} = X_{0:L}X_{L:\infty}$) [14, 15].

The following lemma states a helpful relation between cryptic order and L-merges.

Lemma 1. *Given an* ϵ*-machine with cryptic order* k*: for* $L \leq k$*, there exists an* $L-$*merge; for* $L > k$*, there exists no* $L-$*merge.*

Proof. *See App. A.*

Each L-merge corresponds with a real, positive-semidefinite contribution to some quantum overlap. By Lemma 1, for a cryptic-order k process there is at least one L-merge at each length $L \in [1, \ldots, k]$ and none beyond k. Therefore, at least one overlap receives a real, positive-semidefinite contribution at each length up until k, where there are no further contributions. This leads to our result for overlap accumulation and saturation in terms of the cryptic order.

Theorem 1. *Given a process with cryptic order* k*, for each* $L \in [0, k]$*, each quantum overlap is a nondecreasing function of* L*:*

$$\langle \eta_i(L+1)|\eta_j(L+1)\rangle \geq \langle \eta_i(L)|\eta_j(L)\rangle \ .$$

Furthermore, for each $L \in [1, k]$*, there exists at least one overlap that is increased, as a result of a corresponding* L*-merge. For all remaining* $L \geq k$*, each overlap takes the constant value* $\langle \eta_i(k)|\eta_j(k)\rangle$.

Proof. *See App. A.*

Evidently, the cryptic order is an important length scale not only for classical processes, but also when building efficient quantum encoders.

As an important corollary, this theorem also establishes the relation between cryptic order of the process and depth of its QPMM:

$$d = k + 1 \ . \tag{9}$$

Thus, we discovered the process property that corresponds to a finite DAG QPMM is finite cryptic order. Moreover, the cryptic order corresponds to a topological feature of the QPMM, the depth d, responsible for saturation of the overlaps.

This leads to rephrasing the truncated form of the overlaps sum:

$$\langle \eta_i(L)|\eta_j(L)\rangle = \langle (\sigma_i, \sigma_j)| \sum_{n=0}^{\min(L,k)} \zeta^n \, |\mathrm{SINK}\rangle \ . \tag{10}$$

This form is advantageous for any process that is finite cryptic order. This, of course, includes all finite Markov-order processes—processes used quite commonly in a variety of disciplines.

Since the quantum-compression advantage $C_q(L)$ is a function of only $\boldsymbol{\pi}$ and quantum overlaps, the preceding development also gives a direct lesson about the $C_q(L)$ saturation.

Corollary 1. $C_q(L)$ *has constant value* $C_q(k)$ *for* $L \geq k$.

Proof. *The entropy of an ensemble of pure signal states* $\{p_i, |\psi_i\rangle\}$ *is a function of only probabilities* p_i *and overlaps* $\{\langle\psi_i|\psi_j\rangle\}$. *The result then follows directly from Thm. 1.*

Having established connections among depth, cryptic order, and saturation, we seem to be done analyzing quantum overlap—at least for the finite-cryptic case. To prepare for going beyond finite horizons, however, we should reflect on the spectral origin of ζ's nilpotency.

A nilpotent matrix, such as ζ in the finite-cryptic case, has only the eigenvalue of zero. This can perhaps most easily be seen if the pair-states are ordered according to their distance from SINK, so that ζ is triangular with only zeros along the diagonal.

However, for finite DAGs with depth $d > 1$, the standard eigenvalue–eigenvector decomposition is insufficient to form a complete basis—the corresponding ζ is necessarily nondiagonalizable due to the geometric multiplic-

ity of the zero eigenvalue being less than its algebraic multiplicity. *Generalized eigenvectors* must be invoked to form a complete basis [16]. Intuitively, this type of nondiagonalizability can be understood as the intrinsic interdependence among pair-states in propagating probability amplitude through a branch of the DAG. When ζ is rendered into Jordan block form via a similarity transformation, the size of the largest Jordan block associated with the zero eigenvalue is called the *index* ν_0 of the zero eigenvalue. It turns out to be equal to the depth for finite DAGs.

Summarizing, the finite-horizon case is characterized by several related features: (i) the QPMM is a DAG (of finite depth), (ii) the depth of the QPMM is one greater than the cryptic order, (iii) the matrix ζ has only the eigenvalue zero, and (iv) the depth is equal to the index of this zero-eigenvalue, meaning that ζ has at least k generalized eigenvectors. More generally, ζ can have other eigenvalues and this corresponds to richer structure that we explore next.

B. Infinite Horizon: Lollipop Process

Now we ask, what happens when the QPMM is not a directed acyclic graph? That is, what happens when it contains *cycles*?

It is clear that the depth d diverges, implying that the cryptic order is infinite. Therefore, the sum in Eq. (6) may no longer be truncated. We also know that infinite-cryptic processes become ubiquitous as ϵ-machine state size increases. Have we lost our calculational efficiencies? No, in fact, there are greater advantages yet to be gained.

We first observe that a QPMM's ζ breaks into two pieces. One has a finite horizon reminiscent of finite cryptic order just analyzed, and the other has an infinite horizon, but is analytically quite tractable, as we now show. More generally, a linear operator A may be decomposed using the *Dunford decomposition* [17] (also known as the Jordan–Chevalley decomposition) into:

$$A = \mathcal{D} + \mathcal{N} \ , \tag{11}$$

where $\mathcal{D}$ is diagonalizable, $\mathcal{N}$ is nilpotent, and $\mathcal{D}$ and $\mathcal{N}$ commute. $\mathcal{N}$ makes the familiar finite-horizon contribution, whereas the new $\mathcal{D}$ term has an infinite horizon: $\mathcal{D}^n \neq 0$, for all $n < \infty$. In the context of infinite cryptic processes, the finite horizon associated with $\mathcal{N}$ is no longer simply related to QPMM depth nor, therefore, the cryptic order which is infinite.

The systematic way to address the new diagonalizable part is via a spectral decomposition [18], where the persistent leaky features of the QPMM state probability evolution are understood as independently acting modes. It is clear that ζ always has a nilpotent component associated with a zero eigenvalue, due to the SINK state. Assuming that the remaining eigenspaces are diagonalizable, the form of the overlaps becomes:

$$\langle \eta_i(L) | \eta_j(L) \rangle = \sum_{\lambda \in \Lambda_\zeta \setminus \{0\}} \frac{1-\lambda^{L+1}}{1-\lambda} \langle (\sigma_i, \sigma_j) | \, \zeta_\lambda \, | \text{SINK} \rangle + \sum_{m=0}^{\min\{L, \nu_0 - 1\}} \langle (\sigma_i, \sigma_j) | \, \zeta^m \zeta_0 \, | \text{SINK} \rangle \ , \tag{12}$$

where Λ_ζ is the set of ζ's eigenvalues, ζ_λ are the projection operators corresponding to each eigenvalue, and ν_0 is the index of the zero eigenvalue, which is the size of its largest Jordan block. We refer to this as the *almost-diagonalizable* case. This case covers all processes with generic parameters. Here, ν_0 is still responsible for the length of the finite-horizon component, but is no longer directly related to QPMM depth or process cryptic order.

Note that in the finite-cryptic order case, the only projector ζ_0 is necessarily the identity. Therefore, Eq. (12) reduces to the previous form in Eq. (10).

The spectral decomposition yields a new level of tractability for the infinite-cryptic case. The infinite-horizon piece makes contributions at all lengths, but in a regular way. This allows for direct computation of its total contribution at any particular L, including $L \to \infty$.

To highlight this behavior, consider the (7–4)-Lollipop Process, whose ϵ-machine is shown in Fig. 4. It is named for the shape of its QPMM; see Fig. 5. This process is a simple example of one where the cryptic order is infinite and the finite-horizon length of the nilpotent contribution is tunable. Roughly speaking, the diagonalizable component comes from the "head" of the lollipop (the cycle), and the nilpotent part comes from the "stick".

It is straightforward to construct the general QPMM and thereby derive ζ for the (N–M)-Lollipop Process. Its QPMM has N pair-states in a cyclic head. The M

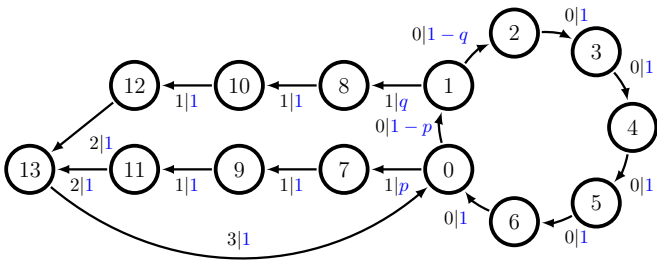

FIG. 4. ϵ-Machine for the (7–4)-Lollipop Process. The cycle of 0s leads to infinite Markov and cryptic orders.

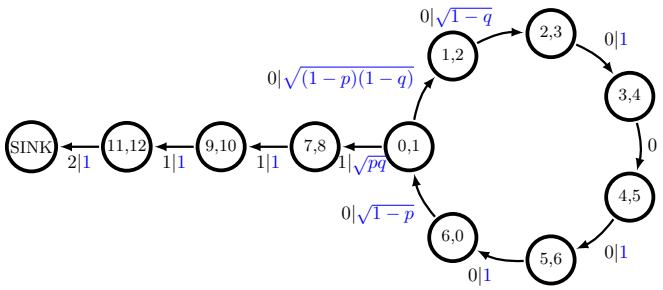

FIG. 5. QPMM for the (7–4)-Lollipop Process.

remaining pair-states constitute a finite-horizon 'stick'. We find:

$$\det(\zeta - \lambda I) = (-\lambda)^M \left[(-\lambda)^N - (1-p)(1-q)\right] ,$$

yielding:

$$\Lambda_\zeta = \left\{0,\ \left[(1-p)(1-q)\right]^{1/N} e^{in2\pi/N}\right\}_{n=0}^{N-1} , \qquad (13)$$

with $\nu_0 = M$.

For concreteness, consider the (7–4)-Lollipop Process with $p = q = 1/2$. It has eigenvalues $\Lambda_\zeta = \{0, ae^{in\theta}\}$ and $\nu_0 = 4$, where $a = (1/4)^{1/7}$, $\theta = 2\pi/7$, and $n \in \{0, 1, 2, 3, 4, 5, 6\}$.

Each $\lambda = ae^{in\theta}$ eigenvalue has algebraic multiplicity 1 and associated left eigenvector:

$$[2\sqrt{2}\lambda^6, \sqrt{2}\lambda^5, \lambda^4, \lambda^3, \lambda^2, \lambda^1, \lambda^0, \sqrt{2}\lambda^5, \sqrt{2}\lambda^4, \sqrt{2}\lambda^3, \sqrt{2}\lambda^2]$$

and right eigenvector:

$$[\frac{1}{2\lambda}, 1, \sqrt{2}\lambda, \sqrt{2}\lambda^2, \sqrt{2}\lambda^3, \sqrt{2}\lambda^4, \sqrt{2}\lambda^5, 0, 0, 0, 0]^\top .$$

Notice that, since ζ is not Hermitian, the right eigenvectors are not simply the conjugate transpose of their left counterparts. The left and right eigenvectors are fundamentally different, with the differences expressing the QPMM's directed causal architecture.

Since each of these eigenvalues has algebraic multiplicity 1, the corresponding projection operators are defined in terms of right and left eigenvectors:

$$\zeta_\lambda = \frac{|\lambda\rangle \langle\lambda|}{\langle\lambda|\lambda\rangle} .$$

The zero eigenvalue has algebraic multiplicity $\nu_0 = 4$ and geometric multiplicity 1, meaning that while there is only one eigenvector there are three generalized eigenvectors. The left and right eigenvectors are:

$$[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1] \text{ and}$$
$$[0, 1, 0, 0, 0, 0, 0, -1, 0, 0, 0]^\top .$$

The three generalized left eigenvectors are:

$$[0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0] ,$$
$$[0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0] , \text{ and}$$
$$[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0] ;$$

and the three generalized right eigenvectors are:

$$[0, 0, 2, 0, 0, 0, 0, 0, -1, 0, 0]^\top ,$$
$$[0, 0, 0, 2, 0, 0, 0, 0, 0, -1, 0]^\top , \text{ and}$$
$$[0, 0, 0, 0, 2, 0, 0, 0, 0, 0, -1]^\top .$$

Since the index of the zero eigenvalue is larger than 1 ($\nu_0 = 4$), the projection operator ζ_0 for the zero eigenvalue includes the contributions from both its standard and generalized eigenvectors:

$$\zeta_0 = \sum_{n=0}^{3} \frac{|\lambda_n\rangle \langle\lambda_n|}{\langle\lambda_n|\lambda_n\rangle} , \qquad (14)$$

where $|\lambda_0\rangle$ is the standard eigenvector and $|\lambda_n\rangle$ is the n^{th} generalized eigenvector for $n \geq 1$. More generally, this sum goes over all standard and all generalized eigenvectors of the zero eigenvalue.

Since all projection operators must sum to the identity, the projection operator for the zero eigenvalue can be obtained alternatively from:

$$\zeta_0 = I - \sum_{\lambda \in \Lambda_\zeta \setminus 0} \zeta_\lambda , \qquad (15)$$

which is often a useful observation during calculations.

Using these projectors, we apply Eq. (12) to compute the overlaps at any L for the (4–7)-Lollipop Process. This very efficient procedure allows us to easily probe the form of quantum advantage for any process described by a finite ϵ-machine.

Finally, we jump directly to the asymptotic overlap using the following expression:

$$\begin{aligned}\langle\eta_i(\infty)|\eta_j(\infty)\rangle &= \langle(\sigma_i, \sigma_j)| \left(\sum_{n=0}^{\infty} \zeta^n\right) |\text{SINK}\rangle \\ &= \langle(\sigma_i, \sigma_j)| \left(I - \zeta\right)^{-1} |\text{SINK}\rangle \ . \qquad (16)\end{aligned}$$

Note that $I - \zeta$ is invertible, since ζ is substochastic. Hence, its spectral radius is less than unity, $1 \notin \Lambda_\zeta$, and so $\det(I - 1\zeta) \neq 0$. Moreover, $(I - \zeta)^{-1}$ is *equal* to the

convergent Neumann series $\sum_{n=0}^{\infty} \zeta^n$ by Thm. 3 of Ref. [19, Ch. VIII § 2].

Yielding an important calculational efficiency, the form of Eq. (16) does not require spectrally decomposing ζ and so immediately provides the asymptotic advantage of quantum compression. Finally, this form does not depend on the previous assumption of ζ being almost-diagonalizable.

VI. QUANTUM COMMUNICATION COST

The preceding development focused on computing overlaps between quantum signal states for q-machine representations of a given process. Let's not forget that the original goal was to compute the von Neumann entropy of this ensemble—the quantum communication cost $C_q(L)$.

The naive approach to calculating $C_q(L)$ constructs the signal states directly and so does not make use of overlap computation. This involves working with a Hilbert space of increasing dimension, exponential in codeword length L. This quickly becomes intractable, for all but the simplest processes.

The second approach, introduced in Ref. [9], made use of the PMM to compute overlaps. These overlaps were then used to construct a density operator with those same overlaps, but in a Hilbert space of fixed size $|\boldsymbol{\mathcal{S}}|$, essentially obviating the high-dimensional embedding of the naive approach. The elements of the resulting density matrix, however, are nonlinear functions of the overlaps. Besides the computational burden this entails, it makes it difficult to use the overlap matrix to theoretically infer much about the general behavior of $C_q(L)$.

Here, we present two markedly improved approaches that circumvent these barriers. We are ultimately interested in the von Neumann entropy which depends only on the spectrum of the density operator. It has been pointed out that the Gram matrix of an ensemble shares the same spectrum [20]. The *Gram matrix* for our ensemble of pure quantum signal states is:

$$G = \begin{bmatrix} \sqrt{\pi_1 \pi_1} \langle \eta_1 | \eta_1 \rangle & \cdots & \sqrt{\pi_1 \pi_{|\boldsymbol{\mathcal{S}}|}} \langle \eta_1 | \eta_{|\boldsymbol{\mathcal{S}}|} \rangle \\ \vdots & \ddots & \vdots \\ \sqrt{\pi_{|\boldsymbol{\mathcal{S}}|} \pi_1} \langle \eta_{|\boldsymbol{\mathcal{S}}|} | \eta_1 \rangle & \cdots & \sqrt{\pi_1 \pi_{|\boldsymbol{\mathcal{S}}|}} \langle \eta_{|\boldsymbol{\mathcal{S}}|} | \eta_{|\boldsymbol{\mathcal{S}}|} \rangle \end{bmatrix} . \tag{17}$$

If we define $D_{\boldsymbol{\pi}} := \mathrm{diag}(\boldsymbol{\pi})$, then $G = D_{\boldsymbol{\pi}}^{1/2} A A^\dagger D_{\boldsymbol{\pi}}^{1/2}$.

Given that it is only a small step from the overlap matrix $AA^\dagger$ to the Gram matrix G, we see the usefulness of the thorough-going overlap analysis above. The spectrum is then computed using standard methods, symbolically or numerically, for G.

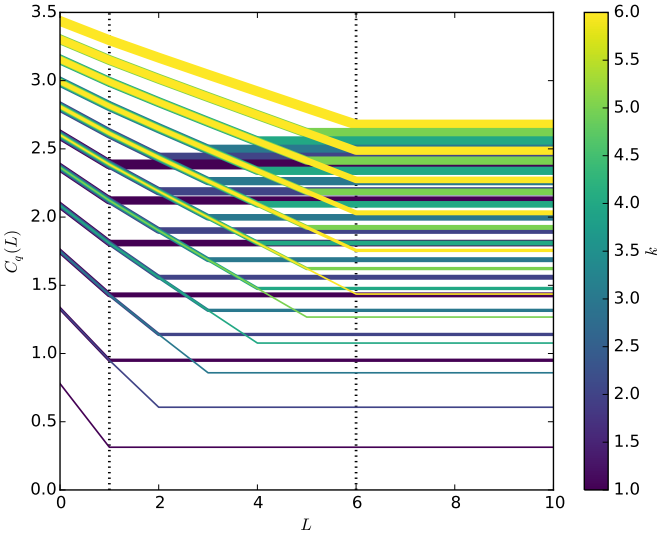

FIG. 6. Quantum costs $C_q(L)$ for the (R–k)-Golden Mean Process with $R \in \{1, \ldots, 6\}$ and $k \in \{1, \ldots, R\}$. R and k are indicated with line width and color, respectively. The probability of the self-loop is $p = 0.7$. $C_q(L)$ roughly linearly decreases until $L = k$ where it is then constant. Note that (R–k) agrees with (($R+1$)–($k-1$)) for $L \leq k$, as explained in App. C.

There is yet another surrogate matrix that also shares the spectrum but is simpler, yet again. We call this matrix $\widetilde{G}$ the *abridged-Gram matrix*:

$$\widetilde{G} = \begin{bmatrix} \pi_1 \langle \eta_1 | \eta_1 \rangle & \cdots & \pi_1 \langle \eta_1 | \eta_{|\boldsymbol{\mathcal{S}}|} \rangle \\ \vdots & \ddots & \vdots \\ \pi_{|\boldsymbol{\mathcal{S}}|} \langle \eta_{|\boldsymbol{\mathcal{S}}|} | \eta_1 \rangle & \cdots & \pi_{|\boldsymbol{\mathcal{S}}|} \langle \eta_{|\boldsymbol{\mathcal{S}}|} | \eta_{|\boldsymbol{\mathcal{S}}|} \rangle \end{bmatrix} . \tag{18}$$

Note that $\widetilde{G} = D_{\boldsymbol{\pi}} A A^\dagger$.

Since the spectra are identical, we can calculate $C_q(L)$ directly from the density matrix $\rho(L)$, Gram matrix G, or abridged-Gram matrix $\widetilde{G}$:

$$\begin{aligned} C_q(L) &= -\sum_{\lambda \in \Lambda_{\rho(L)}} \lambda \log \lambda \\ &= -\sum_{\lambda \in \Lambda_{G(L)}} \lambda \log \lambda \\ &= -\sum_{\lambda \in \Lambda_{\widetilde{G}(L)}} \lambda \log \lambda \ . \end{aligned}$$

For further discussion, see App. B.

Using the Gram matrix as described, we illustrate the behavior of $C_q(L)$ for the (R–k)-Golden Mean (Fig. 6) and (N–M)-Lollipop (Fig. 7). For each of the two process classes, we compute several instances by varying R and k and by varying N and M while holding fixed their transition parameters. Comparing the two figures, we qualitatively confirm the difference between a process with only

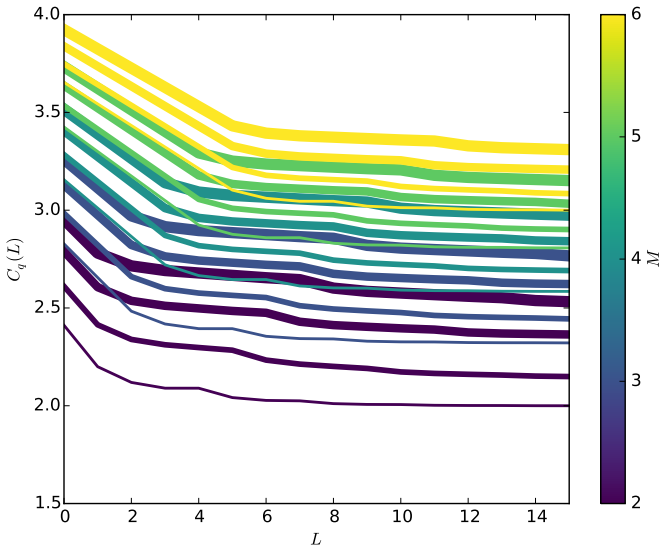

FIG. 7. Quantum costs $C_q(L)$ for the Lollipop process for $N \in \{3,4,5,6\}$, $M \in \{2,3,4,5,6\}$, and $p = q = 0.5$. N and M are indicated with line width and color, respectively. After a fast initial decrease, these curves approach their asymptotic values more slowly.

a finite-horizon contribution and one with an infinite-horizon contribution. The $(R$–$k)$-Golden Mean reaches its encoding saturation at $L = k$ the cryptic order. The $(N$–$M)$-Lollipop only approaches this limit asymptotically.

To supplement the details already given, annotated analytic derivations of several example processes are given in App. C. These examples serve as a pedagogical resource, with comparison and discussion of various analytical techniques.

VII. COSTS USING LONG FINITE CODEWORDS

We discussed quantum-state overlaps extensively. We found that the behavior of the overlaps with L is completely described through ζ's spectral decomposition. And, we showed that, for any L, the von Neumann entropy $C_q(L)$ can be found from the eigenvalues of the Gram matrix—a direct transformation of the overlap matrix. This is all well and good, and key progress. But, can we use this machinery to directly analyze the behavior of $C_q(L)$ as a function of L? For infinite-cryptic processes, the answer is an especially pleasing affirmative.

This section derives $C_q(L)$'s asymptotic behavior for large L; viz., $\nu_0 < L \leq k = \infty$. We show that a periodic pattern, exponentially decaying at the rate of the largest ζ-eigenvalue magnitude, dominates the deviation from $C_q(\infty)$ for large L. That is, we show two things: First, the asymptotic behavior of $C_q(L) - C_q(\infty)$ is, to first order, exponentially decreasing as r_1^L, where r_1 is the spectral radius of ζ. Second, this exponential defines an envelope for a Ψ-periodic asymptotic structure, where Ψ is the least common multiple of slowest-decaying QPMM cycle lengths.

Recall that the optimal quantum compression is given by the asymptotic von Neumann entropy:

$$C_q(\infty) = -\sum_{\lambda^{(\infty)} \in \Lambda_{G^{(\infty)}}} \lambda^{(\infty)} \log\left(\lambda^{(\infty)}\right) .$$

We will show that when L is large, $(\delta G)^{(L)} := G^{(L)} - G^{(\infty)}$ can be treated as a perturbation to $G^{(\infty)}$. From the corresponding small variations $\left\{(\delta\lambda)^{(L)}\right\}_{\lambda \in \Lambda_G}$, direct computation of the first differential yields the approximate change in the von Neumann entropy:

$$(\delta S)^{(L)} = -\sum_{\lambda \in \Lambda_G} \left[\log\left(\lambda^{(\infty)}\right) + 1\right] (\delta\lambda)^{(L)} , \qquad (19)$$

so long as no zero eigenvalues of $G^{(\infty)}$ vanish at finite L. Our task, therefore, is to find $(\delta\lambda)^{(L)}$ from $(\delta G)^{(L)}$ in terms of ζ's spectral properties.

For easy reference, we first highlight our notation:

- $G^{(L)}$ is a Gram matrix at length L corresponding to $\rho(L)$.
- $\lambda^{(L)} \in \Lambda_{G^{(L)}}$ is any one of its eigenvalues.
- $|\lambda^{(L)}\rangle$ and $\langle\lambda^{(L)}|$ are the right and left eigenvectors of $G^{(L)}$ corresponding to $\lambda^{(L)}$, respectively.
- $(\delta G)^{(L)} := G^{(L)} - G^{(\infty)}$ is the perturbation to $G^{(\infty)}$ investigated here.
- $\xi \in \Lambda_\zeta$ is an eigenvalue of the QPMM transition dynamic ζ.

If using G's symmetric version, the right and left eigenvectors are simply transposes of each other: $\langle\lambda^{(L)}| = \left(|\lambda^{(L)}\rangle\right)^\top$. For simplicity of the proofs, we assume nondegeneracy of $G^{(L)}$'s eigenvalues, so that the projection operator associated with $\lambda^{(L)}$ is $|\lambda^{(L)}\rangle\langle\lambda^{(L)}| \,/\, \langle\lambda^{(L)}|\lambda^{(L)}\rangle$, where the denominator assures normalization. Nevertheless, the eigenbasis of $G^{(L)}$ is always complete, and the final result Thm. 3 retains general validity.

Here, we show that the matrix elements of $(\delta G)^{(L)}$ are arbitrarily small for large enough L, such that first-order perturbation is appropriate for large L, and give the exact form of $(\delta G)^{(L)}$ for use in the computation of $(\delta\lambda)^{(L)}$.

Proposition 1. *For $L \geq \nu_0$, the exact change in Gram matrix is:*

$$(\delta G)^{(L)} = -\sum_{\xi \in \Lambda_\zeta \setminus 0} \frac{\xi^{L+1}}{1-\xi} C_\xi ,$$

where C_ξ is independent of L and has matrix elements:

$$(C_\xi)_{i,j} = \sqrt{\pi_i \pi_j}\, \langle (\sigma_i, \sigma_j) | \zeta_\xi | \mathrm{SINK} \rangle \ .$$

Proof. *We calculate:*

$$\begin{aligned}(\delta G)^{(L)}_{i,j} &= G^{(L)}_{i,j} - G^{(\infty)}_{i,j} \\ &= \sqrt{\pi_i \pi_j} \left(\langle \eta^{(L)}_i | \eta^{(L)}_j \rangle - \langle \eta^{(\infty)}_i | \eta^{(\infty)}_j \rangle \right) \\ &= -\sqrt{\pi_i \pi_j}\, \langle (\sigma_i, \sigma_j) | \zeta^{L+1} (1-\zeta)^{-1} | \mathrm{SINK} \rangle \ .\end{aligned}$$

If we assume that all nonzero eigenvalues of ζ correspond to diagonalizable subspaces, then for $L \geq \nu_0$, the elements of $(\delta G)^{(L)}$ have the spectral decomposition:

$$(\delta G)^{(L)}_{i,j} = -\sum_{\xi \in \Lambda_\zeta \setminus 0} \frac{\xi^{L+1}}{1-\xi} \sqrt{\pi_i \pi_j}\, \langle (\sigma_i, \sigma_j) | \zeta_\xi | \mathrm{SINK} \rangle \ .$$

Since this decomposition is common to all matrix elements, we can factor out the $\left\{ \frac{\xi^{L+1}}{1-\xi} \right\}_\xi$, leaving the L-independent set of matrices:

$$\left\{ C_\xi : (C_\xi)_{i,j} = \sqrt{\pi_i \pi_j}\, \langle (\sigma_i, \sigma_j) | \zeta_\xi | \mathrm{SINK} \rangle \right\}_{\xi \in \Lambda_\zeta} \ ,$$

such that:

$$(\delta G)^{(L)} = -\sum_{\xi \in \Lambda_\zeta \setminus 0} \frac{\xi^{L+1}}{1-\xi} C_\xi \ .$$

Proposition 2. *At large L, the first-order correction to $\lambda^{(\infty)}$ is:*

$$(\delta \lambda)^{(L)} = -\sum_{\xi \in \Lambda_\zeta \setminus 0} \frac{\xi^{L+1}}{1-\xi} \frac{\langle \lambda^{(\infty)} | \, C_\xi \, | \lambda^{(\infty)} \rangle}{\langle \lambda^{(\infty)} | \lambda^{(\infty)} \rangle} \ . \tag{20}$$

Proof. *Perturbing $G^{(\infty)}$ to $G^{(\infty)} + (\delta G)^{(L)}$, the first order change in its eigenvalues is given by:*

$$(\delta \lambda)^{(L)} = \frac{\langle \lambda^{(\infty)} | \, (\delta G)^{(L)} \, | \lambda^{(\infty)} \rangle}{\langle \lambda^{(\infty)} | \lambda^{(\infty)} \rangle} \ , \tag{21}$$

which is standard first-order nondegenerate perturbation theory familiar in quantum mechanics, with the allowance for unnormalized bras and kets. Proposition 2 then follows directly from Eq. (21) *and Prop. 1.*

Theorem 2. *At large L, such that $\nu_0 < L \leq k = \infty$, the first-order correction to $C_q(\infty)$ is:*

$$\begin{aligned}C_q(L) - C_q(\infty) &\approx (\delta S)^{(L)} \\ &= \sum_{\xi \in \Lambda_\zeta \setminus 0} \frac{\xi^{L+1}}{1-\xi} \sum_{\lambda^{(\infty)} \in \Lambda_{G^{(\infty)}}} \langle C_\xi \rangle \left[\log(\lambda^{(\infty)}) + 1 \right] \ , \end{aligned} \tag{22}$$

where:

$$\langle C_\xi \rangle := \frac{\langle \lambda^{(\infty)} | \, C_\xi \, | \lambda^{(\infty)} \rangle}{\langle \lambda^{(\infty)} | \lambda^{(\infty)} \rangle} \ .$$

Proof. *This follows directly from Eq.* (19) *and Prop. 2.*

The large-L behavior of $C_q(L) - C_q(\infty)$ is a sum of decaying complex exponentials. And, to first order, we can even calculate the coefficient of each of these contributions.

Notice that the only L-dependence in Prop. 2 and Thm. 2 come in the form of exponentiating eigenvalues of the QPMM transition dynamic ζ. For very large L, the dominant structure implied by Prop. 2 and Thm. 2 can be teased out by looking at the relative contributions from ζ's first- and second-largest magnitude sets of eigenvalues.

Let r_1 be the spectral radius of ζ, shared by the largest eigenvalues $\Lambda(r_1)$: $r_1 := \max_{\xi \in \Lambda_\zeta} |\xi|$. And, let $\Lambda(r_1) := \mathrm{argmax}_{\xi \in \Lambda_\zeta} |\xi|$. Then, let r_2 be the second-largest magnitude of all of ζ's eigenvalues that differs from r_1: $r_2 := \max_{\xi \in \Lambda_\zeta \setminus \Lambda(r_1)} |\xi|$. And, let $\Lambda(r_2) := \mathrm{argmax}_{\xi \in \Lambda_\zeta \setminus \Lambda(r_1)} |\xi|$. Multiple eigenvalues can belong to $\Lambda(r_1)$. Similarly, multiple eigenvalues can belong to $\Lambda(r_2)$.

Then, $0 \leq (r_2/r_1) < 1$ if ζ has at least one nonzero eigenvalue. This is the case of interest here since we are addressing those infinite-horizon processes with $k = \infty > \nu_0$. Hence, as L becomes large, $(r_2/r_1)^L$ vanishes exponentially if it is not already zero. This leads to a corollary of Prop. 2.

Corollary 2. *For $L \geq \nu_0$, the leading deviation from $\lambda^{(\infty)}$ is:*

$$(\delta \lambda)^{(L)} = -r_1^{L+1} \sum_{\xi \in \Lambda(r_1)} \frac{(\xi/|\xi|)^{L+1}}{1-\xi} \langle C_\xi \rangle \left[1 + O\Big(\big(\tfrac{r_2}{r_1}\big)^L \Big) \right] \ .$$

Notice that $\xi/|\xi|$ lies on the unit circle in the complex plane. Due to their origin in cyclic graph structure, we expect each $\xi \in \Lambda(r_1)$ to have a phase in the complex plane that is a rational fraction of 2π. Hence, there is some n for which $(\xi/|\xi|)^n = 1$, for all $\xi \in \Lambda(r_1)$. The minimal such n, call it Ψ, will be of special importance:

$$\Psi := \min \left\{ n \in \mathbb{N} : (\xi/|\xi|)^n = 1 \text{ for all } \xi \in \Lambda(r_1) \right\} \ . \tag{23}$$

Since all $\xi \in \Lambda(r_1)$ originate from cycles in ζ's graph, we have the result that Ψ is equal to the least common multiple of the cycle lengths implicated in $\Lambda(r_1)$.

For example, if all $\xi \in \Lambda(r_1)$ come from the same cycle in the graph of ζ, then $\Psi = |\Lambda(r_1)|$ and:

$$\Lambda(r_1) = \left\{\xi_m = r_1 e^{im2\pi/|\Lambda(r_1)|}\right\}_{m=1}^{|\Lambda(r_1)|} .$$

That is, $\left\{\xi_m/|\xi_m|\right\}_{m=1}^{|\Lambda(r_1)|}$ are the $|\Lambda(r_1)|^{\text{th}}$ roots of unity, uniformly distributed along the unit circle. If, however, $\Lambda(r_1)$ comes from multiple cycles in ζ's graph, then the least common multiple of the cycle lengths should be used in place of $|\Lambda(r_1)|$.

Recognizing the Ψ-periodic structure of $(\xi/|\xi|)^n$ yields a more informative corollary of Prop. 2:

Corollary 3. *For $L \geq \nu_0$, the leading deviation from $\lambda^{(\infty)}$ is:*

$$(\delta\lambda)^{(L)} = -r_1^{L+1} \sum_{\xi \in \Lambda(r_1)} \frac{(\xi/|\xi|)^{\bmod (L+1, \Psi)}}{1-\xi} \langle C_\xi \rangle \times \left[1 + O\left(\left(r_2/r_1\right)^L\right)\right] .$$

Hence:

$$(\delta\lambda)^{(L+\Psi)} \approx r_1^\Psi \, (\delta\lambda)^{(L)} . \tag{24}$$

We conclude that asymptotically a pattern—of changes in the density-matrix eigenvalues (with period Ψ)—decays exponentially with decay rate of r_1^Ψ per period. There are immediate implications for the pattern of asymptotic changes in $C_q(L)$ at large L.

Corollary 4. *For $L \geq \nu_0$, the leading deviation from $C_q(\infty)$ is:*

$$\begin{aligned} C_q(L) - C_q(\infty) &\approx (\delta S)^{(L)} \\ &= r_1^{L+1} \sum_{\xi \in \Lambda(r_1)} \frac{(\xi/|\xi|)^{\bmod (L+1, \Psi)}}{1-\xi} \\ &\quad \times \sum_{\lambda^{(\infty)} \in \Lambda_{G^{(\infty)}}} \langle C_\xi \rangle \log(\lambda^{(\infty)}) \left[1 + O\left(\left(r_2/r_1\right)^L\right)\right] . \end{aligned}$$

The most profound implication of this detailed analysis can be summarized succinctly.

Theorem 3. *For sufficiently large L:*

$$\frac{C_q(L+\Psi) - C_q(\infty)}{C_q(L) - C_q(\infty)} \approx r_1^\Psi . \tag{25}$$

That is, asymptotically a pattern—of changes in $C_q(L) - C_q(\infty)$ (with period Ψ)—decays exponentially with decay rate of r_1^Ψ per period [21].

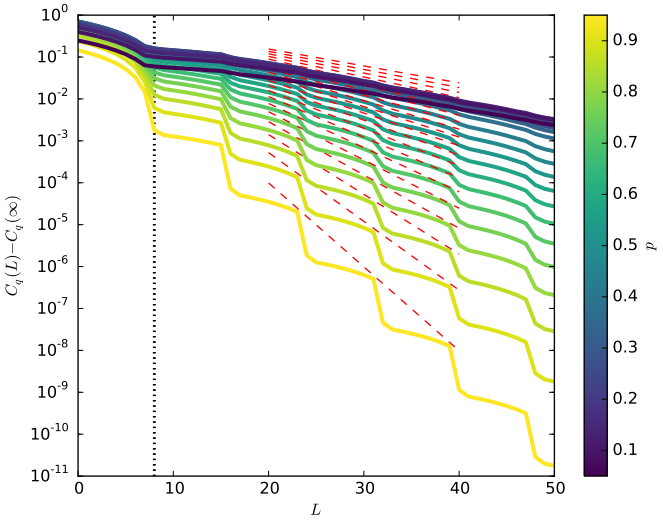

FIG. 8. (8,8)-Lollipop with $p \in [0.05, 0.95]$. $C_q(L) - C_q(\infty)$ on semilog plot illustrates asymptotically exponential behavior. Red dashed lines, r_1^L where r_1 is the spectral radius of ζ, quantify the exponential rate of decay.

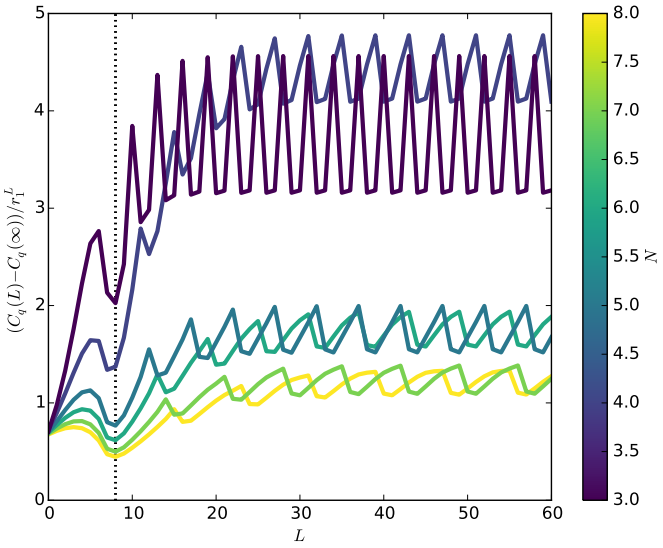

FIG. 9. Lollipop with $N \in \{3, 4, 5, 6, 7, 8\}$ and $M = 8$. $\left(C_q(L) - C_q(\infty)\right)/r_1^L$ demonstrates the periodicity of asymptotic behavior. Removing the exponential envelope makes periodicity of the remaining deviation more apparent. For Lollipop, the periodicity $\Psi = |\Lambda(r_1)| = N$.

While the first-order perturbation allowed us to identify both the roles and values of r_1 and Ψ for any process and Coro. 4 would imply Thm. 3, Thm. 3 actually transcends the limitations of the first-order approximation.

Proof. *Expanding* $\log G^{(L)}$ *in powers of* $(G^{(L)} - I)$, *then multiplying by* $-G^{(L)}$, *shows that* $C_q(L) =$

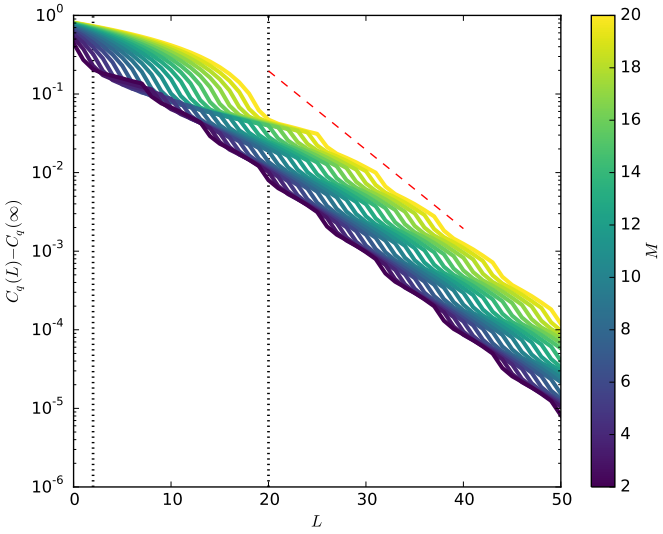

FIG. 10. $C_q(L) - C_q(\infty)$ for Lollipop with $N = 6$ and $M \in \{2, \ldots, 20\}$ on a semilog plot. M determines the finite-horizon length, where the nilpotent part of ζ vanishes. Vertical lines indicate $L = 2$ and $L = 20$, the shortest and longest such length in this group.

$-tr\left[G^{(L)} \log G^{(L)}\right]$ *can be written as:*

$$C_q(L) = -\sum_{n=0}^{\infty} a_n \, tr\left[(G^{(L)})^n\right] \ , \tag{26}$$

for proper $a_n \in \mathbb{R}$. *Using:*

$$G^{(L)} = \sum_{\xi \in \Lambda_\zeta \setminus 0} \frac{1 - \xi^{L+1}}{1 - \xi} C_\xi + \sum_{m=0}^{\min\{L, \nu_0 - 1\}} C_{0,m} \ , \tag{27}$$

with appropriate constant matrices $C_{0,m}$, *together with Eqs.* (23) *and* (26) *yields Thm. 3 with general validity.*

In the simplest case, when ζ has only one largest eigenvalue, then $\Psi = |\Lambda(r_1)| = 1$ and so $C_q(L) - C_q(\infty)$ is dominated by a simple exponential decay at large L.

For the case of multiple largest eigenvalues originating from the same cycle in ζ's graph, then $\Psi = |\Lambda(r_1)| > 1$. And so, the asymptotic behavior of $C_q(L) - C_q(\infty)$ is dominated by a decaying *pattern* of length $|\Lambda(r_1)|$.

For example, the Lollipop processes have an exponentially decaying pattern of length-N that dominates $C_q(L) - C_q(\infty)$ for $L > \nu_0 = M$:

$$\Psi = |\Lambda(r_1)| = N \ . \tag{28}$$

This periodic behavior is apparent in the semi-log plots of Figs. 8 and 10 and is especially emphasized in Fig. 9 which shows that $\Psi = N$ for various N. The figures demonstrate excellent agreement with our qualitative expectations from the above approximations.

Showing the effect of different ν_0, Fig. 10 emphasizes that the initial rolloff of $C_q(L) - C_q(\infty)$ is due to $L \leq \nu_0 = M$. The dominant asymptotic behavior is reached soon after $L = \nu_0$ in this case since the remaining (i.e., nonzero) eigenvalues of the QPMM transition dynamic ζ are all in the largest-magnitude set $\Lambda(r_1)$. In other words, Thm. 2's Eq. (25) is not only approximated by but, in this case, also equal to the simpler expression in Coro. 4, since $r_2 = 0$.

The slope r_1 indicated in Figs. 8 and 10 corresponds to the asymptotic decay rate of the envelope for $C_q(L) - C_q(\infty)$. This asymptotic decay rate is a function of both N and p, since for Lollipop:

$$r_1 = \left[(1-p)(1-q)\right]^{1/N} \ . \tag{29}$$

Figure 8 shows that we have indeed identified the correct slope for different p.

The central asymptotic features of $C_q(L) - C_q(\infty)$ are all captured succinctly by Thm. 3: First, the asymptotic behavior of $C_q(L) - C_q(\infty)$ is exponentially decreasing at a rate of r_1, which is the spectral radius of ζ. Second, this exponential envelope is modulated by an asymptotic Ψ-periodic structure, where Ψ is the least common multiple of slowest-decaying QPMM cycle-lengths.

VIII. CONCLUSION

We developed a detailed analytical theory of how to optimally compress a classical, stationary finite-memory stochastic process using a quantum channel. This required introducing a new quantum state-machine representation (q-machines), carefully constructing its codewords and quantitatively monitoring their overlaps (via the quantum pairwise-merger machine), and utilizing a new matrix formulation of the overlap density matrix (abridged Gram matrix). Applying spectral decomposition then lead directly to closed-form expressions for the quantum coding costs at any codeword length, including infinite length.

The theoretical advances give an extremely efficient way to probe the behavior of quantum compression, both analytically and, when symbolic calculation become arduous, numerically. Analyzing selected example processes illustrated the required calculations and also the range of phenomena to be expected when compressing memoryful processes.

Particular phenomena we reported for the first time here included (i) details of how a process's cryptic order determines its quantum compressibility, (ii) transient and persistent contributions to quantum compression costs, (iii) exponential convergence to optimum compression,

and (iv) oscillations in the convergence that reveal how a process gives up its crypticity with increasing codeword length. Results that hold for finite and infinite Markov and cryptic order processes.

The overall result appears as a rather complete quantitative theory of the informational properties of quantum channels used to compress classical processes, including finite and infinite codeword closed-form expressions. That said, many issues remain, both technical and philosophical. We believe, however, that the approach's mathematical grounding and analytical and numerical efficiency will go some distance to solving them in the near future.

For example, one of the abiding questions is the meaning of process crypticity $\chi = C_\mu - \mathbf{E}$—the difference between a process' predictable information or excess entropy $\mathbf{E}$ and its stored state information or statistical complexity C_μ [10, 22]. Most directly, χ measures how much state information (C_μ) is hidden from observation ($\mathbf{E}$). Cryptic processes and even those with infinite cryptic order dominate the space of classical processes [12]. This means that generically we can compress C_μ down to $C_q(L)$. However, this begs the question of what crypticity is in the quantum domain. Now that we can work analytically in the infinite-length limit, we can explore the *quantum crypticity* $\chi_q = C_q(\infty) - \mathbf{E}$. From our studies, some not reported here, it appears that one cannot compress the state information all the way down to the excess entropy. Why? Why do not quantum models exist of "size" $\mathbf{E}$ bits? Does this point to a future, even more parsimonious physical theory? Or, to a fundamental limitation of communication that even nature must endure, as it channels the past through the present to the future?

For another, are we really justified in comparing Shannon bits (C_μ) to qubits (C_q)? This is certainly not a new or recent puzzle. However, the results on compression bring it to the fore anew. And, whatever the outcome, the answer will change our view of what physical pattern and structure are. Likely, the answer will have a profound effect. Assuming the comparison is valid, why is there a level of classical reality that is more structurally complex when, as we demonstrated and now can calculate, processes are more compactly represented quantum mechanically?

ACKNOWLEDGMENTS

We thank Ryan James for helpful conversations. The authors thank the Santa Fe Institute for its hospitality during visits. JPC is an SFI External Faculty member. This material is based upon work supported by, or in part by, the John Templeton Foundation and U. S. Army Research Laboratory and the U. S. Army Research Office under contract W911NF-13-1-0390.

Appendix A: Quantum Overlaps and Cryptic Order

Lemma 1. Given an ϵ-machine with cryptic order k: for $L \leq k$, there exists an $L-$merge; for $L > k$, there exists no $L-$merge.

Proof. *By definition of cryptic order k:*

$$\mathrm{H}[\mathcal{S}_k | X_{0:\infty}] = 0 \ ,$$

which means that for any given $x_{0:}$ there exists a unique σ_k. Since k is the minimum such length, for $L = k-1$ there exists some word $x_{0:\infty}$ that leaves uncertainty in causal state $\mathcal{S}_{k-1}$. Call two of these uncertain $\mathcal{S}_{k-1}$ states A and B ($A \neq B$). Tracing $x_{0:\infty}$ backwards from A and B, we produce two state paths. These state paths must be distinct at each step due to eM unifilarity—if they were not distinct at some step, they would remain so for all states going forward, particularly at $\mathcal{S}_{k-1}$. The next symbol x_k must take A and B to the same next state F or violate the assumption of cryptic order k. These two state paths and the word $x_{0:k}$ and the final state F make up a k-merger, meaning that cryptic order k implies the existence of a k-merger.

By removing states from the left side of this k-merger, it is easy to see that a k-merger implies the existence of all shorter L-mergers.

By unifilarity again, $\mathrm{H}[\mathcal{S}_k | X_{0:}] = 0 \rightarrow \mathrm{H}[\mathcal{S}_L | X_{0:}] = 0$, for all $L \geq k$. Assume there exists an L-merger for $L > k$ with word w. By definition of L-merger, there is then uncertainty in the state $\mathcal{S}_{L-1}$. This uncertainty exists for any word with w as the prefix—a set with nonzero probability. This contradicts the definition of cryptic order.

Theorem 1. Given a process with cryptic order k, for each $L \in [0, k]$, each quantum overlap $\langle \eta_i(L) | \eta_j(L) \rangle$ is a nondecreasing function of L. Furthermore, for each $L \in [1, k]$, there exists at least one overlap that is increased (as a result of a corresponding L-merge). For all remaining $L \geq k$, each overlap takes a constant value $\langle \eta_i(k) | \eta_j(k) \rangle$.

Proof. *We directly calculate:*

$$\begin{aligned}
\langle \eta_a(L) | \eta_b(L) \rangle &= \sum_{\substack{w, w' \in \mathcal{A}^L \\ j_L, l_L \in \{i\}_{i=1}^M}} \sqrt{T^{(w)}_{a l_L}} \sqrt{T^{(w')}_{b j_L}} \, \langle w | w' \rangle \, \langle \sigma_{l_L} | \sigma_{j_L} \rangle \\
&= \sum_{w, j_L} \sqrt{T^{(w)}_{a j_L}} \sqrt{T^{(w)}_{b j_L}} .
\end{aligned}$$

So, we have:

$$\begin{aligned}
&\langle \eta_a(L+1)|\eta_b(L+1)\rangle \\
&\quad = \sum_{\substack{w'\in\mathcal{A}^{L+1}\\ j_{L+1}}} \sqrt{T^{(w')}_{aj_{L+1}}}\sqrt{T^{(w')}_{bj_{L+1}}} \\
&\quad = \sum_{\substack{w\in\mathcal{A}^{L},s\in\mathcal{A}\\ j_L,l_L,j_{L+1}}} \sqrt{T^{(w)}_{aj_n}}\sqrt{T^{(s)}_{j_nj_{L+1}}}\sqrt{T^{(w)}_{bl_L}}\sqrt{T^{(s)}_{l_Lj_{L+1}}} \\
&\quad = \sum_{\substack{w\in\mathcal{A}^{L},s\in\mathcal{A}\\ j_L,j_{L+1}}} \sqrt{T^{(w)}_{aj_L}}\sqrt{T^{(s)}_{j_Lj_{L+1}}}\sqrt{T^{(w)}_{bj_L}}\sqrt{T^{(s)}_{j_Lj_{L+1}}} \\
&\qquad + \sum_{\substack{w\in\mathcal{A}^{L},s\in\mathcal{A}\\ j_L\neq l_L,j_{L+1}}} \sqrt{T^{(w)}_{aj_L}}\sqrt{T^{(s)}_{j_Lj_{L+1}}}\sqrt{T^{(w)}_{bl_L}}\sqrt{T^{(s)}_{l_Lj_{L+1}}} \ ,
\end{aligned}$$

The first sum represents the overlaps obtained already at length L. To see this, we split the sum to two parts, where the first contains:

$$\begin{aligned}
&\sum_{\substack{w\in\mathcal{A}^{L},s\in\mathcal{A}\\ j_L,j_{L+1}}} \sqrt{T^{(w)}_{aj_L}}\sqrt{T^{(s)}_{j_Lj_{L+1}}}\sqrt{T^{(w)}_{bj_L}}\sqrt{T^{(s)}_{j_Lj_{L+1}}} \\
&\qquad = \sum_{\substack{w\in\mathcal{A}^{L}\\ j_L}} \sqrt{T^{(w)}_{aj_L}}\sqrt{T^{(w)}_{bj_L}}\Big(\sum_{\substack{s\in\mathcal{A}\\ j_{L+1}}} \sqrt{T^{(s)}_{j_Lj_{L+1}}}\sqrt{T^{(s)}_{j_Lj_{L+1}}}\Big) \\
&\qquad = \sum_{\substack{w\in\mathcal{A}^{L}\\ j_L}} \sqrt{T^{(w)}_{aj_L}}\sqrt{T^{(w)}_{bj_L}} \\
&\qquad = \langle \eta_a(L)|\eta_b(L)\rangle \ .
\end{aligned}$$

We use Lemma 1 to analyze the second sum, which represents the change in the overlaps, finding that:

$$\sum_{\substack{w\in\mathcal{A}^{L},s\in\mathcal{A}\\ j_L\neq l_L,j_{L+1}}} \sqrt{T^{(w)}_{aj_L}}\sqrt{T^{(s)}_{j_Lj_{L+1}}}\sqrt{T^{(w)}_{bl_L}}\sqrt{T^{(s)}_{l_Lj_{L+1}}} \geq 0 \ ,$$

with equality when $L \geq k$. Summarizing:

$$\langle \eta_a(L+1)|\eta_b(L+1)\rangle \geq \langle \eta_a(L)|\eta_b(L)\rangle \ ,$$

with equality for $L \geq k$.

Note that while the set of overlaps continues to be augmented at each length up until the cryptic order, we do not currently have a corresponding statement about the nontrivial change in $C_q(L)$ or its monotonicity.

Appendix B: Matrices and Their Entropy

1. Density Matrix

The density matrix can now be expressed using a fixed $|\boldsymbol{\mathcal{S}}|$-by-$|\boldsymbol{\mathcal{S}}|$ matrix, valid for all L. Using the Gram-Schmidt procedure one can choose a new orthonormal basis. Let:

$$\begin{aligned}
|\eta_1(L)\rangle &= |e_1^{(L)}\rangle \\
|\eta_2(L)\rangle &= a_{21}^{(L)}\,|e_1^{(L)}\rangle + a_{22}^{(L)}\,|e_2^{(L)}\rangle \\
|\eta_3(L)\rangle &= a_{31}^{(L)}\,|e_1^{(L)}\rangle + a_{32}^{(L)}\,|e_2^{(L)}\rangle + a_{33}^{(L)}\,|e_3^{(L)}\rangle \\
&\ \ \vdots
\end{aligned}$$

and so on. Then:

$$\begin{aligned}
a_{21}^{(L)} &= \langle \eta_1(L)|\eta_2(L)\rangle \\
&= \langle (\sigma_1,\sigma_2)|\,\big(\sum_{n=0}^{L}\zeta^n\big)\,|\mathrm{SINK}\rangle \ , \\
a_{22}^{(L)} &= \big(1-|\,\langle \eta_1(L)|\eta_2(L)\rangle\,|^2\big)^{1/2} \ , \\
a_{31}^{(L)} &= \langle \eta_1(L)|\eta_3(L)\rangle \\
&= \langle (\sigma_1,\sigma_3)|\,\big(\sum_{n=0}^{L}\zeta^n\big)\,|\mathrm{SINK}\rangle \ ,
\end{aligned}$$

and so on. Now it is useful to rewrite what we can in matrix form:

$$\begin{bmatrix} \langle \eta_1(L)| \\ \langle \eta_2(L)| \\ \langle \eta_3(L)| \\ \vdots \\ \langle \eta_{|\boldsymbol{\mathcal{S}}|}(L)| \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & & & & 0 \\ a_{21}^{(L)} & a_{22}^{(L)} & & & \\ a_{31}^{(L)} & a_{32}^{(L)} & a_{33}^{(L)} & & \\ \vdots & & & \ddots & \\ a_{|\boldsymbol{\mathcal{S}}|1}^{(L)} & & \cdots & & a_{|\boldsymbol{\mathcal{S}}||\boldsymbol{\mathcal{S}}|}^{(L)} \end{bmatrix}}_{\equiv A_L} \begin{bmatrix} \langle e_1^{(L)}| \\ \langle e_2^{(L)}| \\ \langle e_3^{(L)}| \\ \vdots \\ \langle e_{|\boldsymbol{\mathcal{S}}|}^{(L)}| \end{bmatrix} \ ,$$

which defines the lower-triangular matrix A_L. Note that the rightmost matrix of orthonormal basis vectors is simply the identity matrix, since we are working in that basis.

In this new basis, we construct the $|\boldsymbol{\mathcal{S}}|$-by-$|\boldsymbol{\mathcal{S}}|$ density

matrix as:

$$\begin{aligned}\rho(L) &= \sum_{i=1}^{|\boldsymbol{\mathcal{S}}|} \pi_i \left|\eta_i(L)\right\rangle \left\langle\eta_i(L)\right| \\ &= \begin{bmatrix}|\eta_1(L)\rangle & \cdots & |\eta_{|\boldsymbol{\mathcal{S}}|}(L)\rangle\end{bmatrix} \underbrace{\begin{bmatrix}\pi_1 & & 0 \\ & \ddots & \\ 0 & & \pi_{|\boldsymbol{\mathcal{S}}|}\end{bmatrix}}_{\equiv D_\pi} \begin{bmatrix}\langle\eta_1(L)| \\ \langle\eta_2(L)| \\ \langle\eta_3(L)| \\ \vdots \\ \langle\eta_{|\boldsymbol{\mathcal{S}}|}(L)|\end{bmatrix} \\ &= A_L^\dagger D_\pi A_L \ .\end{aligned}$$

Since all entries are real, the conjugate transpose is the transpose. This more general framework may be useful, however, if we want to consider the effect of adding phase to the quantum states.

2. von Neumann Entropy

The quantum coding cost is:

$$\begin{aligned}C_q(L) &= -\mathrm{tr}\left[\rho(L)\log\rho(L)\right] \\ &= -\mathrm{tr}\left[A_L^\dagger D_\pi A_L \log(A_L^\dagger D_\pi A_L)\right] \\ &= -\sum_{\lambda\in\Lambda_{A_L^\dagger D_\pi A_L}} \lambda\log\lambda \ .\end{aligned}$$

This is relatively easy to evaluate since the density matrix $\rho(L)$ is only a $|\boldsymbol{\mathcal{S}}|$-by-$|\boldsymbol{\mathcal{S}}|$ function of L. Thus, we calculate $C_q(L)$ analytically from ρ's spectrum. This, in a curious way, was already folded into ζ's spectrum.

3. Gram Matrix

The A_L matrix is burdensome due to nonlinear dependence on the overlap of the quantum states. We show how to avoid this nonlinearity and instead obtain the von Neumann entropy from a transformation that yields a linear relationship with overlaps.

The *Gram matrix*, with elements $G_{mn}^{(L)} = \sqrt{\pi_m\pi_n}\left\langle\eta_m(L)|\eta_n(L)\right\rangle$, can be used instead of $\rho(L)$ to evaluate the von Neumann entropy [20]. In particular, $G^{(L)}$ has the same spectrum as $\rho(L)$, even with the same multiplicities: $\Lambda_{G^{(L)}} = \Lambda_{\rho(L)}$, while a_λ, g_λ, and ν_λ remain unchanged for all λ in the spectrum. (This is a slightly stronger statement than in Ref. [20], but is justified since $\rho(L)$ and $G^{(L)}$ are both $|\boldsymbol{\mathcal{S}}|$-by-$|\boldsymbol{\mathcal{S}}|$ dimensional.)

Here, we briefly explore the relationship between $\rho(L)$ and $G^{(L)}$ and, then, focus on the closed-form expression for $G^{(L)}$. The result is more elegant than $\rho(L)$, allowing us to calculate and understand $C_q(L)$ more easily.

Earlier, we found that the density matrix can be written as:

$$\rho(L) = A_L^\dagger D_\pi A_L \ ,$$

which can be rewritten as:

$$\begin{aligned}\rho(L) &= A_L^\dagger D_\pi^{1/2} D_\pi^{1/2} A_L \\ &= \left(D_\pi^{1/2} A_L\right)^\dagger D_\pi^{1/2} A_L \ .\end{aligned}$$

It is easy to show that:

$$\begin{aligned}\mathrm{tr}\left[\left(D_\pi^{1/2}A_L\right)^\dagger D_\pi^{1/2}A_L\right] &= \mathrm{tr}\left[D_\pi^{1/2}A_L\left(D_\pi^{1/2}A_L\right)^\dagger\right] \\ &= \mathrm{tr}\left[D_\pi^{1/2}A_L A_L^\dagger D_\pi^{1/2}\right] \ .\end{aligned}$$

This means that the sum of the eigenvalues is conserved in transforming from $A_L^\dagger D_\pi A_L$ to $D_\pi^{1/2}A_L A_L^\dagger D_\pi^{1/2}$. It is less obvious that the spectrum is also conserved, but this is also true, and even easy to prove. (Observe that $AB\vec{v} = \lambda\vec{v} \implies BAB\vec{v} = \lambda B\vec{v} \implies BA(B\vec{v}) = \lambda(B\vec{v})$.) Interestingly, the new object turns out to be exactly the Gram matrix, which was previously introduced, although without this explicit relationship to the density matrix. We now see that:

$$\begin{aligned}&D_\pi^{1/2}A_L A_L^\dagger D_\pi^{1/2} \\ &= D_\pi^{1/2}\begin{bmatrix}\langle\eta_1(L)| \\ \vdots \\ \langle\eta_{|\boldsymbol{\mathcal{S}}|}(L)|\end{bmatrix}\begin{bmatrix}|\eta_1(L)\rangle & \cdots & |\eta_{|\boldsymbol{\mathcal{S}}|}(L)\rangle\end{bmatrix} D_\pi^{1/2} \\ &= \begin{bmatrix}\sqrt{\pi_1}\left\langle\eta_1(L)\right| \\ \vdots \\ \sqrt{\pi_{|\boldsymbol{\mathcal{S}}|}}\left\langle\eta_{|\boldsymbol{\mathcal{S}}|}(L)\right|\end{bmatrix}\begin{bmatrix}\sqrt{\pi_1}\left|\eta_1(L)\right\rangle & \cdots & \sqrt{\pi_{|\boldsymbol{\mathcal{S}}|}}\left|\eta_{|\boldsymbol{\mathcal{S}}|}(L)\right\rangle\end{bmatrix} \\ &= \begin{bmatrix}\sqrt{\pi_1\pi_1}\left\langle\eta_1(L)|\eta_1(L)\right\rangle & \cdots & \sqrt{\pi_1\pi_{|\boldsymbol{\mathcal{S}}|}}\left\langle\eta_1(L)|\eta_{|\boldsymbol{\mathcal{S}}|}(L)\right\rangle \\ \vdots & \ddots & \vdots \\ \sqrt{\pi_{|\boldsymbol{\mathcal{S}}|}\pi_1}\left\langle\eta_{|\boldsymbol{\mathcal{S}}|}(L)|\eta_1(L)\right\rangle & \cdots & \sqrt{\pi_1\pi_{|\boldsymbol{\mathcal{S}}|}}\left\langle\eta_{|\boldsymbol{\mathcal{S}}|}(L)|\eta_{|\boldsymbol{\mathcal{S}}|}(L)\right\rangle\end{bmatrix} \\ &= G^{(L)} \ .\end{aligned}$$

Since the spectrum is preserved, we can use the Gram matrix directly to compute the von Neumann entropy:

$$\begin{aligned}C_q(L) &= -\sum_{\lambda\in\Lambda_{G^{(L)}}} \lambda\log\lambda \\ &= -\mathrm{tr}\left[G^{(L)}\log G^{(L)}\right] \ .\end{aligned}$$

4. Abridged-Gram Matrix

Transforming to the Gram matrix suggests a similar and even more helpful simplification that can be made while preserving the spectrum. Define the *abridged Gram matrix* to be:

$$\begin{aligned}\widetilde{G}^{(L)} &\equiv D_\pi A_L A_L^\dagger \\ &= D_\pi \begin{bmatrix} \langle \eta_1(L)| \\ \vdots \\ \langle \eta_{|\boldsymbol{\mathcal{S}}|}(L)| \end{bmatrix} \begin{bmatrix} |\eta_1(L)\rangle & \cdots & |\eta_{|\boldsymbol{\mathcal{S}}|}(L)\rangle \end{bmatrix} \\ &= \begin{bmatrix} \pi_1 \langle \eta_1(L)|\eta_1(L)\rangle & \cdots & \pi_1 \langle \eta_1(L)|\eta_{|\boldsymbol{\mathcal{S}}|}(L)\rangle \\ \vdots & \ddots & \vdots \\ \pi_{|\boldsymbol{\mathcal{S}}|} \langle \eta_{|\boldsymbol{\mathcal{S}}|}(L)|\eta_1(L)\rangle & \cdots & \pi_{|\boldsymbol{\mathcal{S}}|} \langle \eta_{|\boldsymbol{\mathcal{S}}|}(L)|\eta_{|\boldsymbol{\mathcal{S}}|}(L)\rangle \end{bmatrix}\end{aligned}$$

Clearly, this preserves the same trace as the density matrix and previous Gram matrix. It also preserves the spectrum. And, it has the advantage of not using square-roots of two different state probabilities in each element. Rather it has a single probability attached to each element.

Since the spectrum is preserved, we can use the abridged-Gram matrix to compute the von Neumann entropy:

$$C_q(L) = -\sum_{\lambda \in \Lambda_{\widetilde{G}^{(L)}}} \lambda \log \lambda \tag{B1}$$

$$= -\mathrm{tr}\left[\widetilde{G}^{(L)} \log \widetilde{G}^{(L)}\right] . \tag{B2}$$

Appendix C: Examples

Exploring several more examples will help to illustrate the methods and lead to additional observations.

1. Biased Coins Process

The Biased Coins Process provides a first, simple case that realizes a nontrivial quantum state entropy [7]. There are two biased coins, named A and B. The first generates 1 with probability q; the second, 0 with probability p. A coin is picked and flipped, generating outputs 0 or 1. With probability q the other coin is used next similarly with different probability. Its two causal-state ϵ-machine is shown in Fig. 11.

After constructing the QPMM for the Biased Coins Process, as outlined in Figs. 11 and 12, we observe:

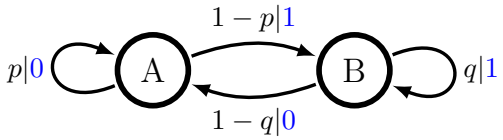

FIG. 11. ϵ-Machine for the Biased Coins Process.

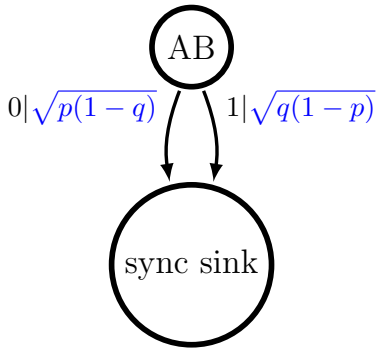

FIG. 12. QPMM for the Biased Coins Process.

$$\zeta^{(0)} = \begin{bmatrix} 0 & \sqrt{p(1-q)} \\ 0 & 0 \end{bmatrix} ,$$

$$\zeta^{(1)} = \begin{bmatrix} 0 & \sqrt{q(1-p)} \\ 0 & 0 \end{bmatrix} ,$$

and so:

$$\zeta = \begin{bmatrix} 0 & \beta \\ 0 & 0 \end{bmatrix} ,$$

where we defined $\beta \equiv \sqrt{p(1-q)} + \sqrt{q(1-p)}$. Let's also define the suggestive quantity $\gamma \equiv (1-\beta^2)^{-1/2}$.

The only overlap to consider is $\langle \eta_A(L)|\eta_B(L)\rangle$. For this, we note that $\langle (A,B)| = \begin{bmatrix} 1 & 0 \end{bmatrix}$. Also, $|\mathrm{SINK}\rangle = \begin{bmatrix} 0 & 1 \end{bmatrix}^\top$.

Spectrally, ζ here is a nilpotent matrix with only a zero eigenvalue with index two: $\Lambda_\zeta = \{0\}$ and $\nu_0 = 2$. Since the projection operators must sum to the identity, we have $\zeta_0 = I$.

ζ^L is the null matrix for $L > 1$, so either by Eq. (8) or by Eq. (10), we have:

$$\langle \eta_A(L)|\eta_B(L)\rangle = \sum_{m=1}^{\min\{L,1\}} \langle (A,B)|\, \zeta^m \,|\mathrm{SINK}\rangle .$$

That is:

$$\langle \eta_A(L)|\eta_B(L)\rangle = \begin{cases} 0 & \text{if } L = 0 , \\ \beta & \text{if } L \geq 1 . \end{cases}$$

a. Entropy from the Density Matrix

For the density matrix, we turn to the L-dependent orthonormal basis $\{|e_1^{(L)}\rangle, |e_2^{(L)}\rangle\}$ and use the stationary distribution over $\boldsymbol{\mathcal{S}}$: $\pi = \left[p/(p+q) \;\; q/(p+q)\right]$.

Apparently, for $L = 0$ we have: $|\eta_A(0)\rangle = |e_1^{(0)}\rangle$ and $|\eta_B(0)\rangle = |e_2^{(0)}\rangle$. Hence, $\rho(0) = D_\pi$ and $C_q(0) = H_2(p/(p+q)) = C_\mu$ qubits.

For $L \geq 1$ we have: $|\eta_A(L)\rangle = |e_1^{(L)}\rangle$ and $|\eta_A(L)\rangle = a_{21}^{(L)}\,|e_1^{(L)}\rangle + a_{22}^{(L)}\,|e_2^{(L)}\rangle$, where $a_{21}^{(L)} = \langle\eta_A(L)|\eta_B(L)\rangle = \beta$ and $a_{22}^{(L)} = (1-\beta^2)^{1/2} = \gamma^{-1}$ for $L \geq 1$. We find that:

$$A_L = \begin{bmatrix} 1 & 0 \\ \beta & \gamma^{-1} \end{bmatrix} \text{ , for } L \geq 1 \text{ .}$$

Hence, for $L \geq 1$ the density matrix is:

$$\begin{aligned} \rho(L) &= A_L^\dagger D_\pi A_L \\ &= \begin{bmatrix} 1 & \beta \\ 0 & \gamma^{-1} \end{bmatrix} \begin{bmatrix} \frac{p}{p+q} & 0 \\ 0 & \frac{q}{p+q} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ \beta & \gamma^{-1} \end{bmatrix} \\ &= \tfrac{1}{p+q} \begin{bmatrix} p & q\beta \\ 0 & q\gamma^{-1} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ \beta & \gamma^{-1} \end{bmatrix} \\ &= \tfrac{q}{p+q} \begin{bmatrix} \frac{p}{q} + \beta^2 & \beta/\gamma \\ \beta/\gamma & 1-\beta^2 \end{bmatrix} \text{ .} \end{aligned}$$

Since:

$$\det(\rho(L) - \lambda I) = \lambda^2 - \lambda + \tfrac{pq}{(p+q)^2}(1-\beta^2) \text{ ,}$$

we find $\rho(L)$'s eigenvalues to be:

$$\Lambda_{\rho(L)} = \left\{ \tfrac{1}{2} \pm \tfrac{1}{2(p+q)} \sqrt{4pq\beta^2 + (p-q)^2} \right\} \text{ ,}$$

which yields the von Neumann entropy for $L \geq 1$:

$$C_q(L) = -\sum_{\lambda \in \Lambda_{\rho(L)}} \lambda \log \lambda \text{ .}$$

b. Entropy from the Abridged-Gram Matrix

The abridged-Gram matrix for the Biased Coins Process is:

$$\widetilde{G}^{(L)} = D_\pi \begin{bmatrix} \langle\eta_A(L)|\eta_A(L)\rangle & \langle\eta_A(L)|\eta_B(L)\rangle \\ \langle\eta_B(L)|\eta_A(L)\rangle & \langle\eta_B(L)|\eta_B(L)\rangle \end{bmatrix} \text{ .}$$

FIG. 13. ϵ-Machine for the $(R$–$k)$-Golden Mean Process.

Specifically, we have for $L \geq 1$:

$$\begin{aligned} \widetilde{G}^{(0)} &= \tfrac{1}{p+q} \begin{bmatrix} p & 0 \\ 0 & q \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \\ &= \tfrac{1}{p+q} \begin{bmatrix} p & 0 \\ 0 & q \end{bmatrix} \end{aligned}$$

and:

$$\begin{aligned} \widetilde{G}^{(L)} &= \tfrac{1}{p+q} \begin{bmatrix} p & 0 \\ 0 & q \end{bmatrix} \begin{bmatrix} 1 & \beta \\ \beta & 1 \end{bmatrix} \\ &= \tfrac{1}{p+q} \begin{bmatrix} p & p\beta \\ q\beta & q \end{bmatrix} \text{ .} \end{aligned}$$

$\widetilde{G}^{(0)}$'s eigenvalues are simply its diagonal entries. So, $C_q(0) = H_2(p/(p+q))$ qubits. For $L \geq 1$:

$$\det(\widetilde{G}^{(L)} - \lambda I) = \lambda^2 - \lambda + \tfrac{pq}{(p+q)^2}(1-\beta^2) \text{ ,}$$

which gives the same value for eigenvalues and entropy.

As the new method illustrates, there is no need to construct the density matrix. Instead, one uses the abridged-Gram matrix, which can be easily calculated from quantum overlaps. Clearly, the abridged-Gram matrix method is more elegant for our purposes. This is evident even at $|\boldsymbol{\mathcal{S}}| = 2$. This is even more critical for more complex processes since A_L grows as $|\boldsymbol{\mathcal{S}}|$ grows.

2. $(R$–$k)$-Golden Mean Process

The $(R$–$k)$-Golden Mean Process is constructed to have Markov-order R and cryptic-order k. Its ϵ-machine is shown in Fig. 13. The 0^{th} state σ_0 has probability $\pi_0 = 1/(R + k - p(R + k - 1))$ while all other states σ_i have probability $\pi_i = (1-p)\pi_0$.

Its QPMM is strictly tree-like with depth $d = k+1$ and maximal width k. All edges have a unit weight except for those edges leaving A-paired states. The latter edges, numbering k in total, have an associated weight of $\sqrt{p}$.

The eigenvalues of the abridged-Gram matrix can be

obtained from:

$$\det(\widetilde{G}^{(L)} - \lambda I) = (\pi_1 - \lambda)^{R+k-\min(L,k)-1} \times \begin{vmatrix} \pi_0 - \lambda & \pi_0\sqrt{p} & \cdots & \pi_0\sqrt{p}^{\min(L,k)} \\ \pi_1\sqrt{p} & \pi_1 - \lambda & & \pi_1\sqrt{p}^{\min(L,k)-1} \\ \vdots & & \ddots & \\ \pi_1\sqrt{p}^{\min(L,k)} & & & \pi_1 - \lambda \end{vmatrix} = 0 \ ,$$

which directly yields the von Neumann entropy. Note that although the $C_q(L)$ is *not* actually linear in L, it appears approximately linear.

We observe that $\boldsymbol{\pi}$ is invariant under the simultaneous change of:

$$R' = R + m \ , \text{ while } \ k' = k - m \ , \tag{C1}$$

for any $m \in \mathbb{Z}$. Although we insist on maintaining $R' \geq k' \geq 0$ for preservation of their functional roles. Furthermore, this transformation preserves $\det(\widetilde{G}^{(L)} - \lambda I)$ for $L \leq k$ and $L \leq k'$. Hence $C_q(L)$ is invariant to the simultaneous transformation of Eq. (C1) for $L \leq k$ and $L \leq k'$, which explains our near-linear observation in Fig. 6's caption.

To give an explicit example, let's consider the (4–3)-GM Process of Fig. 1. State A has probability $\pi_A = 1/(7-6p)$ while all other states have probability $\pi_i = (1-p)/(7-6p)$. Let's calculate:

- For $L = 0$:

$$\det(\widetilde{G}^{(0)} - \lambda I) = (\pi_B - \lambda)^6(\pi_A - \lambda) \ ,$$

yielding $\Lambda_{\widetilde{G}^{(0)}} = \{\pi_B, \pi_A\}$ (with $a_{\pi_B} = 6$) and

$$C_q(0) = -6\pi_B \log \pi_B - \pi_A \log \pi_A \ .$$

- For $L = 1$:

$$\det(\widetilde{G}^{(1)} - \lambda I) = (\pi_B - \lambda)^5 \left[\lambda^2 - (\pi_A + \pi_B)\lambda + \pi_A\pi_B(1-p)\right] \ ,$$

yielding $\Lambda_{\widetilde{G}^{(1)}} = \{\pi_B, c_+, c_-\}$ with $c_\pm = \frac{1}{2}(\pi_A + \pi_B) \pm \frac{1}{2}\left[(\pi_A + \pi_B)^2 - 4\pi_A\pi_B(1-p)\right]^{1/2}$ (and with $a_{\pi_B} = 5$), and:

$$C_q(1) = -5\pi_B \log \pi_B - c_+ \log c_+ - c_- \log c_- \ .$$

- For $L = 2$:

$$\det(\widetilde{G}^{(2)} - \lambda I) = (\pi_B - \lambda)^4 \begin{vmatrix} \pi_A - \lambda & \pi_A p^{1/2} & \pi_A p \\ \pi_B p^{1/2} & \pi_B - \lambda & \pi_B p^{1/2} \\ \pi_B p & \pi_B p^{1/2} & \pi_B - \lambda \end{vmatrix} \ .$$

- For $L \geq 3$:

$$\det(\widetilde{G}^{(L)} - \lambda I) = \det(\widetilde{G}^{(3)} - \lambda I) = (\pi_B - \lambda)^3 \begin{vmatrix} \pi_A - \lambda & \pi_A p^{1/2} & \pi_A p & \pi_A p^{3/2} \\ \pi_B p^{1/2} & \pi_B - \lambda & \pi_B p^{1/2} & \pi_B p \\ \pi_B p & \pi_B p^{1/2} & \pi_B - \lambda & \pi_B p^{1/2} \\ \pi_B p^{3/2} & \pi_B p & \pi_B p^{1/2} & \pi_B - \lambda \end{vmatrix} \ .$$

[1] J. P. Crutchfield. Between order and chaos. *Nature Physics*, 8(January):17–24, 2012.

[2] P. Ball. *The Self-Made Tapestry: Pattern Formation in Nature*. Oxford University Press, New York, 1999.

[3] R. Hoyle. *Pattern Formation: An Introduction to Methods*. Cambridge University Press, New York, 2006.

[4] H. Kantz and T. Schreiber. *Nonlinear Time Series Analysis*. Cambridge University Press, Cambridge, United Kingdom, second edition, 2006.

[5] J. P. Crutchfield and K. Young. Inferring statistical complexity. *Phys. Rev. Let.*, 63:105–108, 1989.

[6] C. R. Shalizi and J. P. Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity. *J. Stat. Phys.*, 104:817–879, 2001.

[7] M. Gu, K. Wiesner, E. Rieper, and V. Vedral. Quantum mechanics can reduce the complexity of classical models. *Nature Comm.*, 3(762):1–5, 2012.

[8] P. Gmeiner. Equality conditions for internal entropies of certain classical and quantum models. *arXiv:1108.5303*, 2011.

[9] J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield. Occam's quantum strop: Synchronizing and compressing classical cryptic processes via a quantum channel. 2015. Santa Fe Institute Working Paper 15-08-030; arxiv.org:1508.02760 [quant-ph].

[10] J. P. Crutchfield, C. J. Ellison, and J. R. Mahoney. Time's barbed arrow: Irreversibility, crypticity, and stored information. *Phys. Rev. Lett.*, 103(9):094101, 2009.

[11] R. B. Ash. *Information Theory*. John Wiley and Sons, New York, 1965.

[12] R. G. James, J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield. Many roads to synchrony: Natural time scales and their algorithms. *Physical Review E*, 89:042135, 2014.

[13] J. P. Crutchfield, C. J. Ellison, J. R. Mahoney, and R. G. James. Synchronization and control in intrinsic and designed computation: An information-theoretic analysis of competing models of stochastic computation. *CHAOS*, 20(3):037105, 2010.

[14] J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield. Information accessibility and cryptic processes. *J. Phys. A: Math. Theo.*, 42:362002, 2009.

[15] J. R. Mahoney, C. J. Ellison, R. G. James, and J. P. Crutchfield. How hidden are hidden processes? a primer on crypticity and entropy convergence. *CHAOS*, 21(3):037112, 2011.

[16] J. N. Franklin. *Matrix Theory*. Dover Publications, New York, 2000.

[17] N. Dunford. Spectral operators. *Pacific J. Math.*, 4(3):321–354, 1954.

[18] J. P. Crutchfield, P. M. Riechers, and C. J. Ellison. Exact complexity: Spectral decomposition of intrinsic computation. 2013. Santa Fe Institute Working Paper 13-09-028; arXiv:1309.3792 [cond-mat.stat-mech].

[19] K. Yosida. *Functional Analysis*. Classics in Mathematics. Cambridge University Press, 1995.

[20] R. Jozsa and J. Schlienz. Distinguishability of states and von Neumann entropy. *Phys. Rev. A*, 62:012301, 2000.

[21] In principle, we need to consider two cases: the pattern decays to $C_q(\infty)$ from above or from below. In either case, the decay of the $C_q(L) - C_q(\infty)$ pattern would be exponential. However, it is known that $C_q(L)$ is strictly less than $C_\mu = C_q(0)$ for any L for any noncounifilar process (and equal otherwise). Hence, we expect that $C_q(L)$ always decays from above, as strongly corroborated by extensive numerical exploration.

[22] C. J. Ellison, J. R. Mahoney, and J. P. Crutchfield. Prediction, retrodiction, and the amount of information stored in the present. *J. Stat. Phys.*, 136(6):1005–1034, 2009.